

An Algebraic Approach to Maliciously Secure Private Set Intersection

Satrajit Ghosh¹ and Tobias Nilges²

¹Department of Computer Science, Aarhus University, Denmark

²ITK Engineering GmbH, Germany

Abstract

Private set intersection is an important area of research and has been the focus of many works over the past decades. It describes the problem of finding an intersection between the input sets of at least two parties without revealing anything about the input sets apart from their intersection.

In this paper, we present a new approach to compute the intersection between sets based on a primitive called Oblivious Linear Function Evaluation (OLE). On an abstract level, we use this primitive to efficiently add two polynomials in a randomized way while preserving the roots of the added polynomials. Setting the roots of the input polynomials to be the elements of the input sets, this directly yields an intersection protocol with optimal asymptotic communication complexity $O(m\kappa)$. We highlight that the protocol is information-theoretically secure assuming OLE.

We also present a natural generalization of the 2-party protocol for the fully malicious multi-party case. Our protocol does away with expensive (homomorphic) threshold encryption and zero-knowledge proofs. Instead, we use simple combinatorial techniques to ensure the security. As a result we get a UC-secure protocol with asymptotically optimal communication complexity $O((n^2 + nm)\kappa)$, where n is the number of parties, m is the set size and κ the security parameter. Apart from yielding an asymptotic improvement over previous works, our protocols are also conceptually simple and require only simple field arithmetic.

Along the way we develop tools that might be of independent interest.

Keywords: Private set intersection, oblivious linear function evaluation, multi-party, UC-security

1 Introduction

Private set intersection (PSI) has been the focus of research for decades and describes the following basic problem. Two parties, Alice and Bob, each have a set S_A and S_B , respectively, and want to find the intersection $S_\cap = S_A \cap S_B$ of their sets. This problem is non-trivial if both parties must not learn anything but the intersection. There are numerous applications for PSI from auctions [NPS99] over advertising [PSSZ15] to proximity testing [NTL⁺11].

Supported by the European Union's Horizon 2020 research and innovation programme under grant agreement #669255 (MPCPRO).

Over the years several techniques for two-party PSI have been proposed, which can be roughly placed in four categories: constructions built from specific number-theoretic assumptions [Sha80, MM87, HFH99, DKT10, DT10], using garbled circuits [HEK12, PSSZ15], based on oblivious transfer (OT) [DCW13, PSZ14, PSZ16, OOS17, KKRT16, RR17] and based on oblivious polynomial evaluation (OPE) [FNP04, DMRY09, HN12, Haz15, FHNP16]. There also exists efficient PSI protocols in server-aided model [KMRS14].

Some of these techniques translate to the multi-party setting. The first (passively secure) multi-party PSI (MPSI) protocol was proposed by Freedman et al. [FNP04] based on OPE and later improved in a series of works [KS05, SS08, CJS12] to achieve full malicious security. Recently, Hazay and Venkatasubramanian [HV17] proposed new protocols secure against semi-honest and fully malicious adversaries. They improve upon the communication efficiency of previous works by designating a central party that runs a version of the protocol of [FNP04] with all other parties and aggregates the results.

Given the state of the art, it remains an open problem to construct a protocol with asymptotically optimal communication complexity in the fully malicious multi-party setting. The main reason for this is the use of zero-knowledge proofs and expensive checks in previous works, which incur an asymptotic overhead over passively secure solutions.

In a concurrent and independent work, Kolesnikov et al. [KMP⁺17] presented a new paradigm for solving the problem of MPSI from oblivious programmable pseudorandom functions (OPPRF). Their approach yields very efficient protocols for multi-party PSI, but the construction achieves only passive security against $n - 1$ corruptions. However, their approach to aggregate the intermediate results uses ideas similar to our masking scheme in the multi-party case.

1.1 Our Contribution

We propose a new approach to (multi-party) private set intersection based on oblivious linear function evaluation (OLE). OLE allows two mutually distrusting parties to evaluate a linear function $ax + b$, where the sender knows a and b , and the receiver knows x . Nothing apart from the result $ax + b$ is learned by the receiver, and the sender learns nothing about x . OLE can be instantiated in the OT-hybrid model from a wide range of assumptions with varying communication efficiency, like LPN [ADI⁺17], Quadratic/Composite Residuosity [IPS09] and Noisy Encodings [IPS09, GNN17], or even unconditionally [IPS09].

Our techniques differ significantly from previous works and follow a new paradigm which leads to conceptually simple and very efficient protocols. This results in asymptotic efficiency improvements over previous works in both communication and computational complexity (cf. Table 1). Our approach is particularly efficient if all input sets are of similar size. To showcase the benefits of our overall approach, we also describe how our MPSI protocol can be modified into a threshold MPSI protocol.

Concretely, we achieve the following:

- Two-party PSI with communication complexity $O(m\kappa)$ and computational complexity $O(m \log m)$. The protocol is information-theoretically secure against a malicious adversary in the OLE-hybrid model.
- UC-secure Multi-party PSI in fully malicious setting with communication complexity $O((n^2 + nm)\kappa)$ and computational complexity $O(nm \log m)$ for the central party and $O(m \log m)$ for

Protocol	Tools	Communication	Computation	Corruptions	Security
[KMP ⁺ 17]	OPPRF	$O(nm\kappa)$	$O(n\kappa)$	$n - 1$	semi-honest
[HV17]	THE	$O(nm\kappa)$	$O(nm \log m)$	$n - 1$	semi-honest
[KS05]	THE, ZK	$O(n^2 m^2 \kappa)$	$O(n^2 m + nm^2 \kappa)$	$n - 1$	malicious
[CJS12]	THE, ZK	$O(n^2 m \kappa)$	$O(n^2 m + nm\kappa)$	$t < n/2$	malicious
[HV17]	CRS, THE	$O((n^2 + nm \log m)\kappa)$	$O(m^2)$	$n - 1$	malicious
Ours + [GNN17]	OLE	$O((n^2 + nm)\kappa)$	$O(nm \log m)$	$n - 1$	malicious ¹

Table 1: Comparison of multi-party PSI protocols, where n is the number of parties, m the size of the input set and κ a security parameter. *The computational cost does not distinguish between exponentiations and multiplications.* Some of the protocols perform better if the sizes of the input sets differ significantly, or particular domains for inputs are used. The overhead described here assumes sets of similar size, with κ bit elements.
the other parties.

- A simple extension of the multi-party PSI protocol to threshold PSI, with the same complexity. To the best of our knowledge, this is the first actively secure threshold multi-party PSI protocol.

In comparison to previous works which rely heavily on exponentiations in fields or groups, our protocols require only field addition and multiplication (and OWF in the case of MPSI). We want to emphasize that this efficiency holds *including* the communication and computation cost for the OLE, if the recent instantiation by Ghosh et al. [GNN17] is used, which is based on noisy Reed-Solomon codes. This OLE protocol has a constant communication overhead and therefore does not influence the asymptotic efficiency of our result. Our results may seem surprising in light of the information-theoretic lower bound of $O(n^2 m \kappa)$ in the communication complexity for multi-party PSI in the fully malicious UC setting. We circumvent this lower bound by considering a slightly modified ideal functionality, resulting in a UC-secure solution for multi-party PSI with asymptotically optimal communication overhead. By asymptotically optimal, we mean that our construction matches the optimal bounds in the plain model for $m > n$, even for passive security, where n is the number of parties, m is the size of the sets and κ is the security parameter. All of our protocols work over fields \mathbb{F} that are exponential in the size of the security parameter κ .

We believe that our approach provides an interesting alternative to existing solutions and that the techniques which we developed can find application in other settings as well.

1.2 Technical Overview

Abstractly, we let both parties encode their input set as a polynomial, such that the roots of the polynomials correspond to the inputs. This is a standard technique, but usually the parties then use OPE to obliviously evaluate the polynomials or some form of homomorphic encryption. Instead, we devise an OLE-based construction to add the two polynomials in an oblivious way, which results in an intersection polynomial. Kissner and Song [KS05] also create an intersection polynomial similar to ours, but encrypted under a layer of homomorphic encryption, whereas our technique results in a *plain* intersection polynomial. Since the intersection polynomial already hides everything but the

¹Our protocol is UC-secure in the fully malicious setting.

intersection, one could argue that the layer of encryption in [KS05] incurs *additional overhead* in terms of expensive computations and complex checks.

In our case, both parties simply evaluate the intersection polynomial on their input sets and check if it evaluates to 0. This construction is information-theoretically secure in the OLE-hybrid model and requires only simple field operations. Conceptually, we compute the complete intersection in one step. In comparison to the naive OPE-based approach², our solution directly yields an asymptotic communication improvement in the input size. Another advantage is that our approach generalizes to the multi-party setting.

We start with a detailed overview of our constructions and technical challenges.

Oblivious polynomial addition from OLE. Intuitively, OLE is the generalization of OT to larger fields, i.e. it allows a sender and a receiver to compute a linear function $c(x) = ax + b$, where the sender holds a, b and the receiver inputs x and obtains c . OLE guarantees that the receiver learns nothing about a, b except for the result c , while the sender learns nothing about x .

Based on this primitive, we define and realize a functionality OPA that allows to add two polynomials in such a way that the receiver cannot learn the sender’s input polynomial, while the sender learns nothing about the receiver’s polynomial or the output. We first describe a passively secure protocol. Concretely, assume that the sender has an input polynomial a of degree $2d$, and the receiver has a polynomial b of degree d . The sender additionally draws a uniformly random polynomial r of degree d . Both parties point-wise add and multiply their polynomials, i.e. they evaluate their polynomials over a set of $2d + 1$ distinct points $\alpha_1, \dots, \alpha_{2d+1}$, resulting in $a_i = a(\alpha_i), b_i = b(\alpha_i)$ and $r_i = r(\alpha_i)$ for $i \in [2d + 1]$. Then, for each of $2d + 1$ OLEs, the sender inputs r_i, a_i , while the receiver inputs b_i and thereby obtains $c_i = r_i b_i + a_i$. The receiver interpolates the polynomial c from the $2d + 1$ (α_i, c_i) and outputs it. Since we assume semi-honest behaviour, the functionality is realized by this protocol.

The biggest hurdle in achieving active security for the above protocol lies in ensuring non-zero b and r . In particular, the protocol has to ensure that the inputs b and r are non-zero. Otherwise, e.g. if $b = 0$, the receiver could learn a . One might think that it is sufficient to perform a coin-toss and verify that the output satisfies the supposed relation, i.e. pick a random x , compute $a(x), b(x), r(x)$ and $c(x)$ and everyone checks if $b(x)r(x) + a(x) = c(x)$ and if $b(x), r(x)$ are non-zero. For $r(x) \neq 0$, the check is actually sufficient, because r must have degree at most d , otherwise the reconstruction fails, and only d points of r can be zero ($r = 0$ would require $2d + 1$ zero inputs). For $b \neq 0$, however, just checking for $b(x) \neq 0$ is not sufficient, because at this point, even if the input $b \neq 0$, the receiver can input d zeroes in the OLE, which in combination with the check is sufficient to learn a completely. We resolve this issue by constructing an enhanced OLE functionality which ensures that the receiver input is non-zero. We believe that this primitive is of independent interest and describe it in more detail later in this section.

Two-party PSI from OLE. Let us first describe a straightforward two-party PSI protocol with one-sided output from the above primitive. Let S_A and S_B denote the inputs for Alice and Bob, respectively, where $|S_P| = m$. They compute \mathbf{p}_A and \mathbf{p}_B such that $\mathbf{p}_P(\gamma) = 0$ for $\gamma \in S_P$. If Bob is supposed to get the intersection, Alice picks a uniformly random polynomial \mathbf{r}_A of degree m and inputs $\mathbf{p}_A, \mathbf{r}_A$ into OPA. Bob inputs \mathbf{p}_B , obtains $\mathbf{p}_\cap = \mathbf{p}_A + \mathbf{p}_B \mathbf{r}_A$ and outputs all $\gamma_j \in S_B$ for which $\mathbf{p}_\cap(\gamma_j) = 0$. Obviously, \mathbf{r}_A does not remove any of the roots of \mathbf{p}_B , and therefore all points γ where $\mathbf{p}_B(\gamma) = 0 = \mathbf{p}_A(\gamma)$ remain in \mathbf{p}_\cap .

²Here we mean an OPE is used for each element of the receiver’s input set. This can be circumvented by clever hashing strategies, e.g. [FNP04, HV17].

However, as a stepping stone for multi-party PSI, we are more interested in protocols that provide output to both parties. If we were to use the above protocol and simply announce \mathbf{p}_\cap to Alice, then Alice could learn Bob’s input. Therefore we have to take a slightly different approach. Let \mathbf{u}_A be an additional random polynomial chosen by Alice. Instead of using her own input in the OPA, Alice uses $\mathbf{r}_A, \mathbf{u}_A$, which gives $\mathbf{s}_B = \mathbf{u}_A + \mathbf{p}_B \mathbf{r}_A$ to Bob. Then they run another OPA in the other direction, i.e. Bob inputs $\mathbf{r}_B, \mathbf{u}_B$ and Alice \mathbf{p}_A . Now, both Alice and Bob have a randomized “share” of the intersection, namely \mathbf{s}_A and \mathbf{s}_B , respectively. Adding \mathbf{s}_A and \mathbf{s}_B yields a masked but correct intersection. We still run into the problem that sending either \mathbf{s}_B to Alice or \mathbf{s}_A to Bob allows the respective party to learn the other party’s input. We also have to use additional randomization polynomials $\mathbf{r}'_A, \mathbf{r}'_B$ to ensure privacy of the final result.

Our solution is to simply use the masks \mathbf{u} to enforce the addition of the two shares. Let us fix Alice as the party that combines the result. Bob computes $\mathbf{s}'_B = \mathbf{s}_B - \mathbf{u}_B + \mathbf{p}_B \mathbf{r}'_B$ and sends it to Alice. Alice computes $\mathbf{p}_\cap = \mathbf{s}'_B + \mathbf{s}_A - \mathbf{u}_A + \mathbf{p}_A \mathbf{r}'_A$. This way, the only chance to get rid of the blinding polynomial \mathbf{u}_B is to add both values. But since each input is additionally randomized via the \mathbf{r} polynomials, Alice cannot subtract her own input from the sum. Since the same also holds for Bob, Alice simply sends the result to Bob.

The last step is to check if the values that are sent and the intersection polynomial are consistent. We do this via a simple coin-toss for a random x , and the parties evaluate their inputs on x and can abort if the relation $\mathbf{p}_\cap = \mathbf{p}_B(\mathbf{r}_A + \mathbf{r}'_B) + \mathbf{p}_A(\mathbf{r}'_A + \mathbf{r}_B)$ does not hold, i.e. \mathbf{p}_\cap is computed incorrectly. This type of check enforces semi-honest behaviour, and was used previously e.g. in [BFO12].

A note on the MPSI functionality. We show that by slightly modifying the ideal functionality for multi-party PSI we get better communication efficiency, without compromising the security at all. A formal definition is given in Section 6.1. Typically, it is necessary for the simulator to extract *all* inputs from the malicious parties, input them into the ideal functionality, and then continue the simulation with the obtained ideal intersection. In a fully malicious setting, however, this requires every party to communicate in $O(m\kappa)$ with every other party—otherwise the input is information-theoretically undetermined and cannot be extracted—which results in $O(n^2 m\kappa)$ communication complexity.

The crucial observation here is that in the setting of multi-party PSI, an intermediate intersection between a single malicious party and *all* honest parties is sufficient for simulation. This is due to the fact that inputs by additional malicious parties can only reduce the size of the intersection, and as long as we observe the additional inputs at some point, we can correctly reduce the intersection in the ideal setting before outputting it. On a technical level, we no longer need to extract all malicious inputs right away to provide a correct simulation of the intersection. Therefore, it is not necessary for every party to communicate in $O(m\kappa)$ with every other party. Intuitively, the intermediate intersection corresponds to the case where all malicious parties have the same input. We therefore argue that the security of this modified setting is identical to standard MPSI up to input substitution of the adversary.

Multi-party PSI. The multi-party protocol is a direct generalization of the two-party protocol, with some small adjustments. We consider a network with a star topology, similar to the recent result of [HV17]. One party is set to be the central party, and all other parties (mainly) interact with this central party to compute the result. The main idea here is to delegate most of the work to the central party, which in turn allows to reduce the communication complexity. Since no party is supposed to get any intermediate intersections, we let each party create an additive sharing of their intersection with the central party.

First, consider the following (incorrect) toy example. Let each party P_i execute the two-party PSI as described above with P_0 , up to the point where both parties have shares $\mathbf{s}_{P_0}^i, \mathbf{s}_{P_i}'$. All parties P_i send their shares \mathbf{s}_{P_i}' to P_0 , who adds all polynomials and broadcasts the output. By design of the protocols and the inputs, this yields the intersection of all parties. Further, the communication complexity is in $O(nm\kappa)$, which is optimal. However, this protocol also allows P_0 to learn all intermediate intersections with the other parties, which is not allowed. Previously, all maliciously secure multi-party PSI protocols used threshold encryption to solve this problem, and indeed it might be possible to use a similar approach to ensure active security for the above protocol. For example, a homomorphic threshold encryption would allow to add all these shares homomorphically, without leaking the intermediate intersections. But threshold encryption incurs a significant computational overhead (and increases the complexity of the protocol and its analysis) which we are unwilling to pay.

Instead, we propose a new solution which is conceptually very simple. We add another layer of masking on the shares \mathbf{s}_{P_i} , which forces P_0 to add *all* intermediate shares—at least those of the honest parties. For this we have to ensure that the communication complexity does not increase, so all parties exchange seeds (instead of sending random polynomials directly), which are used in a PRG to mask the intermediate intersections. This technique is somewhat reminiscent of the pseudorandom secret-sharing technique by Cramer et al. [CDI05]. We emphasize that we do not need any public key operations.

Concretely, all parties exchange a random seed and use it to compute a random polynomial in such a way that every pair of parties P_i, P_j holds two polynomials $\mathbf{v}_{ij}, \mathbf{v}_{ji}$ with $\mathbf{v}_{ij} + \mathbf{v}_{ji} = 0$. Then, instead of sending \mathbf{s}_{P_i}' , each party P_i computes $\mathbf{s}_{P_i}'' = \mathbf{s}_{P_i}' + \sum \mathbf{v}_{ij}$ and sends this value. If P_0 obtains this value, it has to add the values \mathbf{s}_{P_i}'' of all parties to remove the masks, otherwise \mathbf{s}_{P_i}'' will be uniformly random.

Finally, to ensure that the central party actually computed the aggregation, we add a check similar to two-party PSI, where the relation, i.e. computing the sum, is verified by evaluating the inputs on a random value x which is obtained by a multi-party coin-toss.

Threshold (M)PSI. First of all, we clarify the term *threshold PSI*. We consider the setting where all parties have m elements as their input, and the output is only revealed if the intersection of the inputs among all parties is bigger than a certain threshold ℓ . In [HOS17] Hallgren et al. defined this notion for two party setting, and finds application whenever two entities are supposed to be matched once a certain threshold is reached, e.g. for dating websites or ride sharing. In contrast, Kissner and Song [KS05] previously defined an *over-threshold PSI* protocol, where an element appears in the intersection if the number of occurrences of this element among the n parties is bigger than the threshold.

We naturally extend the idea of threshold PSI from [HOS17] to a multi-party settings and propose the first actively secure threshold multi-party PSI protocol. On a high level, our solution uses a similar idea to [HOS17], but we use completely different techniques and achieve stronger security and better efficiency. The main idea is to use a robust secret sharing scheme, and the execution of the protocol basically transfers a subset of these shares to the other parties, one share for each element in the intersection. If the intersection is large enough, the parties can reconstruct the shared value.

Concretely, we only modify the input polynomials of each party P_i for the MPSI protocol and add an additional check. Instead of simply setting \mathbf{p}_i such that $\mathbf{p}_i(\gamma_j) = 0$ for all $\gamma_j \in S_i$, we set

$\mathbf{p}_i(\gamma_j) = 1$. Further, for each of the random polynomials $\mathbf{r}_i, \mathbf{r}'_i$ we set $\mathbf{r}_i(\gamma_j) = \rho_j$ and $\mathbf{r}'_i(\gamma_j) = \rho'_j$, where $\rho_1, \dots, \rho_n, \rho'_1, \dots, \rho'_n$ are the shares of two robust (ℓ, n) -secret sharings of random values s_i^0 and s_i^1 , respectively. Now, by computing the intersection polynomial \mathbf{p}_\cap as before, each party obtains exactly $m_\cap = |S_\cap|$ shares, one for each $\gamma_j \in S_i$. If $m_\cap \geq \ell$ then each party can reconstruct $r_\cap = \sum_{i=1}^n (s_i^0 + s_i^1)$, otherwise the intersection remains hidden completely. Given a successful reconstruction, the set S_\cap can be easily identified by using the locations γ_j of valid shares in \mathbf{p}_\cap .

We believe that our techniques can also be applied to [KS05] or follow-up works, but due to the use of threshold homomorphic encryption and zero-knowledge proofs, the efficiency of these protocols would be inferior to our solution.

A New Flavour of OLE. One of the main technical challenges in constructing our protocols is to ensure a non-zero input into the OLE functionality by the receiver. Recall that an OLE computes a linear function $ax + b$. We define an enhanced OLE functionality (cf. Section 3) which ensures that $x \neq 0$, otherwise the output is uniformly random. Our protocol which realises this functionality makes two black-box calls to a normal OLE and is otherwise purely algebraic.

Before we describe the solution, let us start with a simple observation. If the receiver inputs $x = 0$, an OLE returns the value b . Therefore, it is critical that the receiver cannot force the protocol to output b . One way to achieve this is by forcing the receiver to multiply b with some correlated value via an OLE, lets call it \hat{x} . Concretely, we can use an OLE where the receiver inputs \hat{x} and a random s , while the sender inputs b and obtains $\hat{x}b + s$. Now if the sender uses $a + b\hat{x} + s, 0$ as input for another OLE, where the receiver inputs x , the receiver obtains $ax + b\hat{x}x + sx$. Which means that if $\hat{x} = x^{-1}$ then the receiver can extract the correct output. This looks like a step in the right direction, since for $x = 0$ or $\hat{x} = 0$, the output would not be b . On the other hand, the receiver can now force the OLE to output a by choosing $\hat{x} = 0$ and $x = 1$, so maybe we only shifted the problem.

The final trick lies in masking the output such that it is uniform for inconsistent inputs x, \hat{x} . We do this by splitting b into two shares that only add to b if $x \cdot \hat{x} = 1$. The complete protocol looks like this: the receiver plays the sender for one OLE with input x^{-1}, s , and the sender inputs a random u to obtain $t = x^{-1}u + s$. Then the sender plays the sender for the second OLE and inputs $t + a, b - u$, while the receiver inputs x and obtains $c' = (t + a)x + b - u = ux^{-1}x + sx + ax + b - u = ax + b + sx$, from which the receiver can subtract sx to get the result. A cheating receiver with inconsistent input x^*, \hat{x}^* will get $ax + b + u(x^* \hat{x}^* - 1)$ as an output, which is uniform over the choice of u .

1.3 Structure of the Paper

We start with the definition and construction of the enhanced OLE functionality in Section 3. In Section 4 we define an ideal functionality for the addition of polynomials and describe a protocol that realizes this functionality black-box from OLE. Based on this primitive, we first provide a two-party PSI protocol in Section 5 and later a multi-party PSI protocol in Section 6.

2 Preliminaries

We assume $|\mathbb{F}| \in \theta(2^\kappa)$, where κ is a statistical security parameter. Typically, $x \in \mathbb{F}$ denotes a field element, while $\mathbf{p} \in \mathbb{F}[X]$ denotes a polynomial. Let $\mathcal{M}_0(\mathbf{p})$ denote the zero-set for $\mathbf{p} \in \mathbb{F}[X]$, i.e. $\forall x \in \mathcal{M}_0(\mathbf{p}), \mathbf{p}(x) = 0$.

In the proofs, \hat{x} denotes an element either extracted or simulated by the simulator, while x^* denotes an element sent by the adversary.

We slightly abuse notation and denote by $\mathbf{v} = \text{PRG}_d(s)$ the deterministic pseudorandom polynomial of degree d derived from evaluating PRG on seed s .

2.1 Security Model

We prove our protocol in the Universal Composability (UC) framework of Canetti [Can01]. In the framework, security of a protocol is shown by comparing a real protocol π in the real world with an ideal functionality \mathcal{F} in the ideal world. \mathcal{F} is supposed to accurately describe the security requirements of the protocol and is secure per definition. An environment \mathcal{Z} is plugged either to the real protocol or the ideal protocol and has to distinguish the two cases. For this, the environment can corrupt parties. To ensure security, there has to exist a simulator in the ideal world that produces a protocol transcript indistinguishable from the real protocol, even if the environment corrupts a party. We say π UC-realises \mathcal{F} if for all adversaries \mathcal{A} in the real world there exists a simulator \mathcal{S} in the ideal world such that all environments \mathcal{Z} cannot distinguish the transcripts of the parties' outputs.

Oblivious Linear Function Evaluation. Oblivious Linear Function Evaluation (OLE) is a special case of Oblivious Polynomial Evaluation (OPE). In contrast to OPE, only linear functions can be obliviously evaluated. The sender has as input two values $a, b \in \mathbb{F}$ that determine a linear function $f(x) = a \cdot x + b$ over \mathbb{F} , and the receiver gets to obliviously evaluate the linear function on input $x \in \mathbb{F}$. The receiver will learn only $f(x)$, and the sender learns nothing at all. Consider the ideal functionality in Figure 1.

Functionality \mathcal{F}_{OLE}
1. Upon receiving a message (<code>inputS</code> , (a, b)) from the sender with $a, b \in \mathbb{F}$, verify that there is no stored tuple, else ignore that message. Store a and b and send a message (<code>input</code>) to \mathcal{A} .
2. Upon receiving a message (<code>inputR</code> , x) from the receiver with $x \in \mathbb{F}$, verify that there is no stored tuple, else ignore that message. Store x and send a message (<code>input</code>) to \mathcal{A} .
3. Upon receiving a message (<code>deliver</code> , <code>S</code>) from \mathcal{A} , check if both (a, b) and x are stored, else ignore that message. Send (<code>delivered</code>) to the sender.
4. Upon receiving a message (<code>deliver</code> , <code>R</code>) from \mathcal{A} , check if both (a, b) and x are stored, else ignore that message. Set $y = a \cdot x + b$ and send (<code>output</code> , y) to the receiver.

Figure 1: Ideal functionality for oblivious linear function evaluation.

To instantiate our PSI protocol we choose an efficient batch-OLE protocol based on Noisy Reed-Solomon-Codes from [GNN17]. The protocol has a constant communication overhead of 64 field elements per OLE, and require only $O(\log m)$ simple field operations, where κ -bit prime field gives κ -bit security.

2.2 Commitment from \mathcal{F}_{OLE}

Let us briefly sketch how to obtain an efficient multiple commitment protocol from \mathcal{F}_{OLE} . If we simply use the message m and a random field element r as inputs for \mathcal{F}_{OLE} , while the receiver

queries with a random x , we get a UC-secure commitment. Intuitively, the commitment can be simulated because the simulator knows x and can adjust r , and it can be extracted because the simulator learns the message as an input.

The above commitment protocol, however, does not realise $\mathcal{F}_{\text{mCOM}}$. In order to do so, we have to include the id of the sender pid to prevent man-in-the-middle attacks, in particular copying of the commitment. Interestingly, this poses a difficulty in the \mathcal{F}_{OLE} -hybrid setting, since our message space is limited. So either we pick a larger field \mathbb{F}' such that we can embed in $m' \in \mathbb{F}'$ both $m \in \mathbb{F}$ and pid , or we directly construct a commitment which has a slightly larger message space. We take the second approach in Figure 2.




Figure 2: Π_{mCOM} in the \mathcal{F}_{OLE} -hybrid model.

Lemma 1. *The protocol Π_{mCOM} UC-realizes $\mathcal{F}_{\text{mCOM}}$ in the \mathcal{F}_{OLE} -hybrid model.*

Sketch. Corrupted P_i . The simulator against the committing party observes \hat{m}, \hat{m}' and also r_1, r_2 . It aborts if $\gamma = \hat{m}'/\hat{m} \neq \text{pid}$ for P_i . Otherwise, it sends $(\text{commit}, P_i, P_j, \hat{m})$ to $\mathcal{F}_{\text{mCOM}}$.

Let $(\alpha, \beta_1, \beta_2)$ denote the unveil. This simulation is indistinguishable from the real protocol, since the check of the unveil will always fail if $\alpha \neq \hat{m}$. Let σ denote the outcome of the check from \mathcal{A} 's view.

$$\sigma = \alpha x + \beta_1 - q_1 = (\alpha - \hat{m})x + \beta_1 - r_1.$$

Thus, from \mathcal{A} 's view, σ is uniform over the choice of x . We now know that $\alpha = \hat{m}$. If $\gamma \neq \text{pid}$,

$$\sigma' = \hat{m} \cdot \text{pid} y + \beta_2 - q_2 = (\hat{m} \cdot \text{pid} - \hat{m} \cdot \gamma)y + \beta_2 - r_2.$$

In this case, σ' is uniform over the choice of y and the protocol would abort. In conclusion, the simulator extracts the right input and provides an indistinguishable simulation of the real protocol.

Corrupted P_j . The simulator against the receiving party observes the challenges \hat{x}, \hat{y} and simulates the commit phase with a random input ρ using randomness r_1, r_2 . Upon receiving a message (unveil, \hat{m}) in the ideal setting, the simulator sets $\hat{r}_1 = \hat{q}_1 - \hat{m}\hat{x}$ and $\hat{r}_2 = \hat{q}_2 - \hat{m}\text{pid}\hat{x}$. Since q_1 and q_2 are uniform over the choice of r_1 and r_2 , respectively, the simulation is identically distributed. \square

2.3 Non-malleable Commitments

Roughly, the setting of concurrent non-malleable commitments is as follows. An adversary MIM interacts in a left session with polynomially many committers, while simultaneously interacting with receivers in m right sessions. We denote by $\text{MIM}_{\text{COM}}^m(\mathbf{v}, z)$ the distribution of all values committed by MIM in the right sessions, and $\text{Sim}_{\text{COM}}^m(1^\kappa, z)$ the joint distribution of all values committed by the simulator.

Definition 2. A commitment scheme $\{\text{COM.Commit}, \text{COM.Open}\}$ is said to be m -bounded-concurrent non-malleable if for every PPT MIM, there exists a PPT simulator \mathcal{S} such that the following ensembles are computationally indistinguishable:

$$\{\text{MIM}_{\text{COM}}^m(\{v\}, z)\}_{\kappa \in \mathbb{N}, v \in \{0,1\}^\kappa, z \in \{0,1\}^*} \text{ and } \{\text{Sim}_{\text{COM}}^m(1^\kappa, z)\}_{\kappa \in \mathbb{N}, v \in \{0,1\}^\kappa, z \in \{0,1\}^*}$$

2.4 Technical Lemmas

We state several lemmata which are used to show the correctness of our PSI protocols later on.

Lemma 3. Let $\mathbf{p}, \mathbf{q} \in \mathbb{F}[X]$ be non-trivial polynomials. Then,

$$\mathcal{M}_0(\mathbf{p}) \cap \mathcal{M}_0(\mathbf{p} + \mathbf{q}) = \mathcal{M}_0(\mathbf{p}) \cap \mathcal{M}_0(\mathbf{q}) = \mathcal{M}_0(\mathbf{q}) \cap \mathcal{M}_0(\mathbf{p} + \mathbf{q}).$$

This lemma shows that the sum of two polynomials contains the intersection with respect to the zero-sets of both polynomials.

Proof. Let $\mathcal{M}_\cap = \mathcal{M}_0(\mathbf{p}) \cap \mathcal{M}_0(\mathbf{q})$.

“ \supseteq ”: $\forall x \in \mathcal{M}_\cap: \mathbf{p}(x) = \mathbf{q}(x) = 0$. But $\mathbf{p}(x) + \mathbf{q}(x) = 0$, so $x \in \mathcal{M}_0(\mathbf{p} + \mathbf{q})$.

“ \subseteq ”: It remains to show that there is no x such that $x \in \mathcal{M}_0(\mathbf{p}) \cap \mathcal{M}_0(\mathbf{p} + \mathbf{q})$ but $x \notin \mathcal{M}_\cap$, i.e. $\mathcal{M}_0(\mathbf{p}) \cap (\mathcal{M}_0(\mathbf{p} + \mathbf{q}) \setminus \mathcal{M}_\cap) = \emptyset$. Similarly, $\mathcal{M}_0(\mathbf{q}) \cap (\mathcal{M}_0(\mathbf{p} + \mathbf{q}) \setminus \mathcal{M}_\cap) = \emptyset$.

Assume for the sake of contradiction that $\mathcal{M}_0(\mathbf{p}) \cap (\mathcal{M}_0(\mathbf{p} + \mathbf{q}) \setminus \mathcal{M}_\cap) \neq \emptyset$. Let $x \in \mathcal{M}_0(\mathbf{p}) \cap (\mathcal{M}_0(\mathbf{p} + \mathbf{q}) \setminus \mathcal{M}_\cap)$. Then, $\mathbf{p}(x) = 0$, but $\mathbf{q}(x) \neq 0$, otherwise $x \in \mathcal{M}_\cap$. But this means that $\mathbf{p}(x) + \mathbf{q}(x) \neq 0$, i.e. $x \notin \mathcal{M}_0(\mathbf{p} + \mathbf{q})$. This contradicts our assumption, and we get that $\mathcal{M}_0(\mathbf{p}) \cap (\mathcal{M}_0(\mathbf{p} + \mathbf{q}) \setminus \mathcal{M}_\cap) = \emptyset$.

Symmetrically, we get that $\mathcal{M}_0(\mathbf{q}) \cap (\mathcal{M}_0(\mathbf{p} + \mathbf{q}) \setminus \mathcal{M}_\cap) = \emptyset$. The claim follows. \square

Lemma 4. Let $d \in \text{poly}(\log |\mathbb{F}|)$. Let $\mathbf{p} \in \mathbb{F}[X]$, $\deg(\mathbf{p}) = d$ be a fixed but unknown non-trivial polynomial. Further let $\mathbf{q}_1, \dots, \mathbf{q}_l \in \mathbb{F}[X]$ with $\deg(\mathbf{q}_i) \leq d$.

$$\Pr_{\mathbf{p} \in \mathbb{F}[X]} [(\mathcal{M}_0(\mathbf{p}) \cap \mathcal{M}_0(\sum_{i=1}^l \mathbf{q}_i)) \neq (\mathcal{M}_0(\mathbf{p}) \cap \bigcap_{i=1}^l \mathcal{M}_0(\mathbf{q}_i))] \leq \text{negl}(|\mathbb{F}|).$$

This lemma is basically an extension of Lemma 3 and shows that the sum of several polynomials does not create new elements in the intersection unless the supposedly unknown zero-set of \mathbf{p} can be guessed with non-negligible probability.

Proof. We first observe that $\bigcap_{i=1}^l \mathcal{M}_0(\mathbf{q}_i) \subseteq \mathcal{M}_0(\sum_{i=1}^l \mathbf{q}_i)$: it holds that for all $x \in \bigcap_{i=1}^l \mathcal{M}_0(\mathbf{q}_i)$, $\mathbf{q}_i(x) = 0$ for $i \in [l]$. It follows that $\sum_{i=1}^l \mathbf{q}_i(x) = 0$, i.e. $x \in \mathcal{M}_0(\sum_{i=1}^l \mathbf{q}_i)$.

Now, assume for the sake of contradiction that

$$(\mathcal{M}_0(\mathbf{p}) \cap \mathcal{M}_0(\sum_{i=1}^l \mathbf{q}_i)) \neq (\mathcal{M}_0(\mathbf{p}) \cap \bigcap_{i=1}^l \mathcal{M}_0(\mathbf{q}_i))$$

with non-negligible probability ϵ . Let $\mathcal{M} = \mathcal{M}_0(\sum_{i=1}^l \mathbf{q}_i) \setminus \bigcap_{i=1}^l \mathcal{M}_0(\mathbf{q}_i)$.

Then with probability at least ϵ , the set \mathcal{M} is not empty. Further, we can bound $|\mathcal{M}| \leq d$. Pick a random $x \in \mathcal{M}$. It holds that $\Pr[x \in \mathcal{M}_0(\mathbf{p})] \geq \epsilon/d$, which is non-negligible. But since \mathbf{p} is unknown, so is $\mathcal{M}_0(\mathbf{p})$, and the probability that we can find x , so that $x \in \mathcal{M}_0(\mathbf{p})$ is upper bounded by $d/|\mathbb{F}|$ over \mathbf{p} .

This is a contradiction and the claim follows. \square

Lemma 5. *Let $d, d' \in \text{poly}(\log |\mathbb{F}|)$. Let $\mathbf{r} \in \mathbb{F}[X]$, $\deg(\mathbf{r}) = d$ be a uniformly random polynomial. For all non-trivial $\mathbf{p} \in \mathbb{F}[X]$, $\deg(\mathbf{p}) = d'$,*

$$\Pr_{\mathbf{r} \in \mathbb{F}[X]} [(\mathcal{M}_0(\mathbf{r}) \cap \mathcal{M}_0(\mathbf{p})) \neq \emptyset] \leq \text{negl}(|\mathbb{F}|).$$

This lemma establishes that the intersection of a random polynomial with another polynomial is empty except with negligible probability.

Proof. This follows from the fundamental theorem of algebra, which states that a polynomial of degree d evaluates to 0 in a random point only with probability $d/|\mathbb{F}|$.

Since \mathbf{r} (and therefore all $x \in \mathcal{M}_0(\mathbf{r})$) is uniformly random and $|\mathcal{M}_0(\mathbf{r})| = d$, while $|\mathcal{M}_0(\mathbf{p})| = d'$, we get that

$$\Pr[(\mathcal{M}_0(\mathbf{r}) \cap \mathcal{M}_0(\mathbf{p})) \neq \emptyset] \leq dd'/|\mathbb{F}|.$$

\square

Lemma 6. *Let $d \in \text{poly}(\log |\mathbb{F}|)$. Let $\mathbf{p} \in \mathbb{F}[X]$, $\deg(\mathbf{p}) = d$ be a fixed but unknown non-trivial polynomial. Further let $\mathbf{r} \in \mathbb{F}[X]$, $\deg(\mathbf{r}) = d$ be a uniformly random polynomial. For all non-trivial $\mathbf{q}, \mathbf{s} \in \mathbb{F}[X]$ with $\deg(\mathbf{q}) \leq d$ and $\deg(\mathbf{s}) \leq d$,*

$$\Pr_{\mathbf{r} \in \mathbb{F}[X]} [(\mathcal{M}_0(\mathbf{p}) \cap \mathcal{M}_0(\mathbf{ps} + \mathbf{rq})) \neq (\mathcal{M}_0(\mathbf{p}) \cap \mathcal{M}_0(\mathbf{q}))] \leq \text{negl}(|\mathbb{F}|).$$

This lemma shows that the multiplication of (possibly maliciously chosen) polynomials does not affect the intersection except with negligible probability, if one random polynomial is used.

Proof.

$$\begin{aligned} \mathcal{M}_0(\mathbf{p}) \cap \mathcal{M}_0(\mathbf{ps} + \mathbf{rq}) &\stackrel{\text{Lemma 3}}{=} \mathcal{M}_0(\mathbf{p}) \cap (\mathcal{M}_0(\mathbf{ps}) \cap \mathcal{M}_0(\mathbf{qr})) \\ &= \mathcal{M}_0(\mathbf{p}) \cap ((\mathcal{M}_0(\mathbf{p}) \cup \mathcal{M}_0(\mathbf{s})) \cap (\mathcal{M}_0(\mathbf{q}) \cup \mathcal{M}_0(\mathbf{r}))) \\ &= \mathcal{M}_0(\mathbf{p}) \cap ((\mathcal{M}_0(\mathbf{p}) \cap \mathcal{M}_0(\mathbf{q})) \cup \underbrace{(\mathcal{M}_0(\mathbf{p}) \cap \mathcal{M}_0(\mathbf{r}))}_{\mathcal{T}_1}) \\ &\quad \cup \underbrace{(\mathcal{M}_0(\mathbf{s}) \cap \mathcal{M}_0(\mathbf{q}))}_{\subseteq \mathcal{M}_0(\mathbf{q})} \cup \underbrace{(\mathcal{M}_0(\mathbf{s}) \cap \mathcal{M}_0(\mathbf{r}))}_{\mathcal{T}_2} \end{aligned}$$

From Lemma 5 it follows that $\Pr[\mathcal{T}_1 \neq \emptyset] \leq d^2/|\mathbb{F}|$, and also $\Pr[\mathcal{T}_2 \neq \emptyset] \leq d^2/|\mathbb{F}|$. Since

$$\mathcal{M}_0(\mathbf{p}) \cap ((\mathcal{M}_0(\mathbf{p}) \cap \mathcal{M}_0(\mathbf{q})) \cup \mathcal{M}_0(\mathbf{q})) = \mathcal{M}_0(\mathbf{p}) \cap \mathcal{M}_0(\mathbf{q}),$$

we get

$$\Pr_{\mathbf{r} \in \mathbb{F}[X]} [(\mathcal{M}_0(\mathbf{p}) \cap \mathcal{M}_0(\mathbf{p}\mathbf{s} + \mathbf{r}\mathbf{q})) \neq (\mathcal{M}_0(\mathbf{p}) \cap \mathcal{M}_0(\mathbf{q}))] \leq 2d^2/|\mathbb{F}|.$$

□

3 Enhanced Oblivious Linear Function Evaluation $\mathcal{F}_{\text{OLE}^+}$

In this section we present an enhanced version of the OLE functionality. The standard OLE functionality allows the sender to input a, b , while the receiver inputs x and obtains $ax + b$. For our applications, we do not want the receiver to be able to learn b , i.e. it has to hold that $x \neq 0$. Our approach is therefore to modify the OLE functionality in such a way that it outputs a random field element upon receiving an input $x = 0$ (cf. Figure 3). A different approach might be to output a special abort symbol or 0, but crucially the output must *not* satisfy the relation $ax + b$. This is a particularly useful feature, as we will show in the next section.

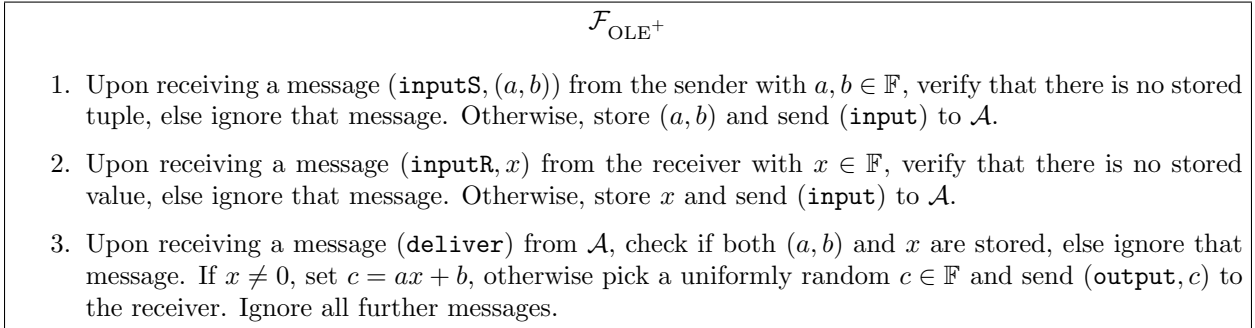


Figure 3: Ideal functionality for the enhanced oblivious linear function evaluation.

While it might be possible to modify existing OLE protocols in such a way that a non-zero input is guaranteed, we instead opt to build a protocol black-box from the standard OLE functionality \mathcal{F}_{OLE} .

Let us begin with a short overview of our construction. Our main goal is to prevent leakage of the value b on input $x = 0$. One way to ensure that is by using an OLE and computing xb . However, the result of the enhanced OLE should still be $ax + b$ for $x \neq 0$. Thus, we have to remove the connection of x and b with a second OLE, while at the same time adding another connection with a . The key idea is to force the receiver to use x^{-1} and x as his inputs, thus cancelling the relation between x and b if the inputs were chosen correctly.

Concretely, one first attempt might be to execute an OLE from receiver to sender, where the receiver inputs x^{-1} and a random value s . The sender inputs b and obtains $bx^{-1} + s$, then adds a and uses $(bx^{-1} + a + s, 0)$ for the second OLE. The receiver inputs x and thereby obtains $ax + b + sx$ from which he can subtract sx (because he knows both s and x) and output $c = ax + b$.

This leaves a small problem: the receiver can still meddle with his inputs, and particularly by first using 0 instead of x^{-1} and then 1 instead of x the protocol returns a . This obviously should

not be allowed. Therefore, we make a small modification to the above described protocol. Instead of sending b to the first OLE, the sender picks a uniformly random u and obtains $ux^{-1} + s$. In order to still get the correct result, the sender's input into the second OLE has to be changed to $(ux^{-1} + a + s, b - u)$. If the receiver cheats with his inputs, the term u will completely randomize the output, since $c = ax + b + uxx^{-1} - u + sx$. The formal description of the protocol is given in Figure 4.




Figure 4: Protocol that realizes $\mathcal{F}_{\text{OLE}^+}$ in the \mathcal{F}_{OLE} -hybrid model.

Lemma 7. Π_{OLE^+} unconditionally UC-realizes $\mathcal{F}_{\text{OLE}^+}$ in the \mathcal{F}_{OLE} -hybrid model.

Proof. The simulator against a corrupted sender simulates both instances of \mathcal{F}_{OLE} . Let α_1 be the sender's input in the first OLE, and (α_2, α_3) be the inputs into the second OLE. The simulator sets $\hat{b} = \alpha_1 + \alpha_3$ and $\hat{a} = \alpha_2 - \hat{t}$, where \hat{t} is chosen as the uniformly random output to \mathcal{A}_S of the first OLE. The simulator simply inputs $(\text{inputS}, (\hat{a}, \hat{b}))$ into $\mathcal{F}_{\text{OLE}^+}$. Let us briefly argue that this simulation is indistinguishable from a real protocol run. The value \hat{t} is indistinguishable from a valid t , since the receiver basically uses a one-time-pad s to mask the multiplication. Therefore, the sender can only change his inputs into the OLEs. Since his inputs uniquely determine both \hat{a} and \hat{b} , the extraction by the simulator is correct and the simulation is indistinguishable from a real protocol run.

Against a corrupted receiver, the simulator simulates the two instance of \mathcal{F}_{OLE} and obtains the receiver's inputs (ξ_1, ξ_3) and ξ_2 . If $\xi_1 \cdot \xi_2 = 1$, the simulator sets $\hat{x} = \xi_2$, sends (inputR, \hat{x}) to $\mathcal{F}_{\text{OLE}^+}$ and receives (output, c) . It forwards $c' = c + \xi_2\xi_3$ to \mathcal{A}_R . If $\xi_1 \cdot \xi_2 \neq 1$, the simulator sends $(\text{inputR}, 0)$ to $\mathcal{F}_{\text{OLE}^+}$ and forwards the output c to the receiver. It remains to argue that this simulation is indistinguishable from the real protocol. From \mathcal{A} 's view, the output c is determined as

$$c = u\xi_1\xi_2 + a\xi_2 + b - u + \xi_2\xi_3 = a\xi_2 + b + u(\xi_1\xi_2 - 1) + \xi_2\xi_3.$$

We can ignore the last term, since it is known to \mathcal{A} . If $\xi_1\xi_2 \neq 1$, then $u(\xi_1\xi_2 - 1)$ does not vanish and the result will be uniform over the choice of u . Thus, by using ξ_2 as the correct input otherwise, we extract the correct value and the simulation is indistinguishable from the real protocol. \square

4 Randomized Polynomial Addition from OLE

Concretely, we have two parties, the sender with a polynomial of degree $2d$ as input and the receiver with a polynomial of degree d as input. The goal is that the receiver obtains the sum of these two polynomials such that it cannot learn the sender's polynomial fully. We want to achieve

this privacy property by using a randomization polynomial that prevents the receiving party from simply subtracting its input from the result. This functionality is defined in Figure 5.

Notice that we have some additional requirements regarding the inputs of the parties. First, the degree of the inputs has to be checked, but the functionality also makes sure that the receiver does not input a 0 polynomial, because otherwise he might learn the input of the sender. Also note that the functionality leaks some information about the sender’s polynomial. Looking ahead in the PSI protocol, where the input of the sender is always a uniformly random $2d$ degree polynomial, this leakage of the ideal functionality will not leak any non-trivial information in the PSI protocol.




Figure 5: Ideal functionality that allows to obliviously compute an addition of polynomials.

4.1 Passively Secure Protocol for \mathcal{F}_{OPA}

It is instructive to first consider a passively secure protocol. In the semi-honest case, both sender and receiver evaluate their input polynomials on a set of distinct points $\mathcal{P} = \{\alpha_1, \dots, \alpha_{2d+1}\}$, where d is the degree of the input polynomials. The sender additionally picks a random polynomial $\mathbf{r} \in \mathbb{F}[X]$ of degree d and also evaluates it on \mathcal{P} .

Instead of using OLE in the “traditional” sense, i.e. instead of computing $\mathbf{a}\mathbf{b} + \mathbf{r}$ where \mathbf{r} blinds the multiplication of the polynomials, we basically compute $\mathbf{r}\mathbf{b} + \mathbf{a}$. This means that the sender randomizes the polynomial of the receiver, and then adds his own polynomial. This prevents the receiver from simply subtracting his input polynomial and learning \mathbf{a} . In a little more detail, sender and receiver use $2d + 1$ OLEs to add the polynomials as follows: for each $i \in [2d + 1]$, the sender inputs (r_i, a_i) in OLE i , while the receiver inputs b_i and obtains $s_i = r_i b_i + a_i$. He then interpolates the resulting polynomial \mathbf{s} of degree $2d$ using the $2d + 1$ values s_i . The above protocol is described in Figure 6.

Due to the security guarantees of OLE, the sender learns nothing about the result, while the receiver’s input is randomized by \mathbf{r} , i.e. he is not able to reconstruct \mathbf{a} .

Lemma 8. $\Pi_{\text{OPA}}^{\text{sh}}$ UC-realizes \mathcal{F}_{OPA} with passive security in the \mathcal{F}_{OLE} -hybrid model.

Sketch. Since both parties provide inputs according to the protocol, it is ensured that the input polynomials have the correct degrees and \mathbf{b} and \mathbf{r} are non-zero. Given the inputs of the parties, we can clearly provide a perfect simulation of the protocol by simulating the OLEs. It remains to argue that $b_i \neq 0$ in $\Pi_{\text{OPA}}^{\text{sh}}$ for honest inputs, since otherwise the receiver can learn the sender’s

$$\Pi_{\text{OPA}}^{\text{sh}}$$

Let $\mathcal{P} = \{\alpha_1, \dots, \alpha_{2d+1}\}, \alpha_i \in \mathbb{F}$ be a set of distinct points.

1. Sender (Input $\mathbf{a} \in \mathbb{F}[X], \deg(\mathbf{a}) = 2d$): Pick $\mathbf{r} \in \mathbb{F}[X]$ of degree d uniformly at random. Evaluate \mathbf{a}, \mathbf{r} on \mathcal{P} to obtain $(a_i, r_i), i \in [2d+1]$. Input (r_i, a_i) into $\mathcal{F}_{\text{OLE}}^{(i)}$.
2. Receiver (Input $\mathbf{b} \in \mathbb{F}[X], \deg(\mathbf{b}) = d$): Evaluate \mathbf{b} on \mathcal{P} and obtain $b_i, i \in [2d+1]$. Input b_i into $\mathcal{F}_{\text{OLE}}^{(i)}$ and receive $s_i = r_i b_i + a_i$. Reconstruct \mathbf{s} from the s_i and output \mathbf{s}

Figure 6: Protocol that realizes \mathcal{F}_{OPA} in the \mathcal{F}_{OLE} -hybrid model with passive security.

input. This follows from the fundamental theorem of algebra: we evaluate \mathbf{b} in $2d+1$ distinct points independent of \mathbf{b} , and \mathbf{b} has exactly d roots. The probability that a root is contained in \mathcal{P} is therefore $(2d+1) \cdot d/|\mathbb{F}|$, which is negligible in $|\mathbb{F}|$. \square

4.2 Actively Secure Protocol for \mathcal{F}_{OPA}

In going from passive to active security, we have to ensure that the inputs of the parties are correct. Here, the main difficulty obviously lies in checking for $\mathbf{b} = 0$. In fact, since \mathcal{F}_{OPA} does not even leak a single point a_i we have to make sure that all $b_i \neq 0$. However, this can easily be achieved by using $\mathcal{F}_{\text{OLE}^+}$ instead of \mathcal{F}_{OLE} . We also have to verify that the inputs are well-formed via a simple polynomial check. For a more detailed overview we refer the reader to the introduction.

The complete actively secure protocol is shown in Figure 7.

$$\Pi_{\text{OPA}}$$

Let $\mathcal{P} = \{\alpha_1, \dots, \alpha_{2d+1}\}, \alpha_i \in \mathbb{F}$ be a set of distinct points.

1. Sender (Input $\mathbf{a}, \mathbf{r} \in \mathbb{F}[X], \deg(\mathbf{a}) \leq 2d, \deg(\mathbf{r}) = d$): Evaluate \mathbf{a}, \mathbf{r} on \mathcal{P} to obtain $(a_i, r_i), i \in [2d+1]$. Input (r_i, a_i) into $\mathcal{F}_{\text{OLE}^+}^{(i)}$.
2. Receiver (Input $\mathbf{b} \in \mathbb{F}[X], \deg(\mathbf{b}) \leq d$): Evaluate \mathbf{b} on \mathcal{P} and obtain $b_i, i \in [2d+1]$. Input b_i into $\mathcal{F}_{\text{OLE}^+}^{(i)}$ and receive $s_i = a_i + b_i \cdot r_i$. Reconstruct \mathbf{s} from the s_i and check if $\deg(\mathbf{s}) = 2d$, otherwise abort.
3. Sender: Pick a random $x_S \in \mathbb{F}$ and send it to the receiver.
4. Receiver: Compute $\mathbf{b}(x_S), \mathbf{s}(x_S)$ and pick a random $x_R \in \mathbb{F}$. Send $(\mathbf{b}(x_S), \mathbf{s}(x_S), x_R)$ to the sender.
5. Sender: If $\mathbf{s}(x_S) \neq \mathbf{a}(x_S) + \mathbf{b}(x_S) \cdot \mathbf{r}(x_S)$, abort. Send $(\mathbf{a}(x_R), \mathbf{r}(x_R))$ to the receiver.
6. Receiver: If $\mathbf{s}(x_R) \neq \mathbf{a}(x_R) + \mathbf{b}(x_R) \cdot \mathbf{r}(x_R)$ or $\mathbf{r}(x_R) = 0$ or $\mathbf{a}(x_R) = 0$, abort, otherwise output \mathbf{s} .

Figure 7: Protocol that realizes \mathcal{F}_{OPA} in the $\mathcal{F}_{\text{OLE}^+}$ -hybrid model.

Lemma 9. Π_{OPA} unconditionally UC-realizes \mathcal{F}_{OPA} in the $\mathcal{F}_{\text{OLE}^+}$ -hybrid model.

Proof. Corrupted Sender. The simulator \mathcal{S}_S against a corrupted sender proceeds as follows. It simulates $\mathcal{F}_{\text{OLE}^+}^{(i)}$ and thereby obtains (r_i^*, a_i^*) for all $i \in [2d+1]$. From these values, the simulator

reconstructs $\hat{\mathbf{r}}$ and $\hat{\mathbf{a}}$. It aborts in Step 6 if $\deg(\hat{\mathbf{r}}) > d$ or $\deg(\hat{\mathbf{a}}) > 2d$. It also aborts if $\hat{\mathbf{a}}$ or $\hat{\mathbf{r}}$ are zero, and otherwise sends $(\text{inputS}, (\hat{\mathbf{a}}, \hat{\mathbf{r}}))$ to \mathcal{F}_{OPA} .

The extraction of the corrupted sender's inputs is correct if his inputs \mathbf{r}^* corresponds to a polynomial of degree d and \mathbf{a}^* corresponds to a polynomial of degree $2d$. Thus, the only possibility for an environment to distinguish between the simulation and the real protocol is by succeeding in answering the check while using a malformed input, i.e. a polynomial of incorrect degree or 0-polynomials. If the polynomials have degree greater than d and $2d$, respectively, the resulting polynomial \mathbf{s} has degree $2d + 1$ instead of $2d$, i.e. the receiver cannot reconstruct the result from $2d + 1$ points. Since the sender learns nothing about the receiver's inputs, the thus incorrectly reconstructed polynomial will be uniformly random from his point of view and the probability that his response to the challenge is correct is $1/|\mathbb{F}|$. Also, both $\hat{\mathbf{a}}$ and $\hat{\mathbf{r}}$ have to be non-zero, because in each case the polynomials are evaluated in $2d + 1$ points, and it requires $2d + 1$ zeros as a_i, r_i to get a 0 polynomial. But since both \mathbf{a}, \mathbf{r} have degree at most $2d$, there are at most $2d$ roots of these polynomials. Therefore, in order to pass the check, $\mathbf{a}(x)$ and $\mathbf{b}(x)$ would need to be 0, which is also checked for.

Corrupted Receiver. The simulator \mathcal{S}_{R} against a corrupted receiver simulates $\mathcal{F}_{\text{OLE}^+}^{(i)}$ and obtains \mathbf{b}_i^* for all $i \in [2d + 1]$. It reconstructs $\hat{\mathbf{b}}$ and aborts the check in Step 5 if $\deg(\hat{\mathbf{b}}) > d$. The simulator sends $(\text{inputR}, \hat{\mathbf{b}})$ to \mathcal{F}_{OPA} and receives $(\text{res}, \hat{\mathbf{s}})$. It evaluates $\hat{\mathbf{s}}$ on \mathcal{P} and returns s_i for the corresponding OLEs. \mathcal{S}_{R} simulates the rest according to the protocol.

Clearly, if the corrupted receiver \mathcal{A}_{R} inputs a degree d polynomial, the simulator will extract the correct polynomial. In order to distinguish the simulation from the real protocol, the adversary can either input 0 in an OLE or has to input a polynomial of higher degree, while still passing the check. In the first case, assume w.l.o.g. that \mathcal{A}_{R} cheats in $\mathcal{F}_{\text{OLE}^+}^{(j)}$ for some j . This means \mathcal{A}_{R} receives a value \hat{s}_i , which is uniformly random. This means that only with probability $1/|\mathbb{F}|$ will \hat{s}_i satisfy the relation $\mathbf{r}\mathbf{b} + \mathbf{a}$ and the check will fail. In the second case, the resulting polynomial would be of degree $2d + 1$, while the receiver only gets $2d + 1$ points of the polynomial. Therefore the real polynomial is underdetermined and \mathcal{A} can only guess the correct value $\hat{\mathbf{s}}(x)$, i.e. the check will fail with overwhelming probability. \square

Remark. We use the abstraction of \mathcal{F}_{OPA} to modularize the construction of our PSI protocols. For practical purposes it is possible to remove the check in Π_{OPA} , since this check only ensures that the inputs and output of the protocol are well-formed polynomials and this can also be checked directly in the PSI protocols.

5 Maliciously Secure Two-party PSI

In this section we provide a maliciously secure two-party PSI protocol with output for *both* parties, i.e. we realize \mathcal{F}_{PSI} as described in Figure 8.

We briefly sketch the protocol in the following; a more detailed overview can be found in the introduction. First, Alice and Bob simply transform their input sets into polynomials. Then, both compute a randomized share of the intersection via our previously defined OPA in such a way that Alice can send her share to Bob without him being able to learn her input. This can be achieved by adding a simple mask to the intermediate share. Bob adds both shares and sends the output to Alice. The protocol only requires two OPA and a simple check which ensures semi-honest

\mathcal{F}_{PSI}
<ol style="list-style-type: none"> 1. Upon receiving a message (<code>input</code>, P, S_P) from party $P \in \{\mathbf{A}, \mathbf{B}\}$, store the set S_P. Once all inputs are given, set $S_\cap = S_\mathbf{A} \cap S_\mathbf{B}$ and send (<code>output</code>, S_\cap) to \mathcal{A}. 2. Upon receiving a message (<code>deliver</code>) from \mathcal{A}, send (<code>output</code>, S_\cap) to the honest party.

Figure 8: Ideal functionality \mathcal{F}_{PSI} for two-party PSI.

behaviour, and no computational primitives. A formal description is given in Figure 9.

$\Pi_{2\text{PSI}}$
Let $m = \max_i S_i + 1$.
Computation of Intersection
<ol style="list-style-type: none"> 1. Alice (Input $S_\mathbf{A}$): Compute a polynomial $\mathbf{p}_\mathbf{A}$ of degree m such that $\mathbf{p}_\mathbf{A}(\gamma_j) = 0$ for all $\gamma_j \in S_\mathbf{A}$. Generate two random polynomials $\mathbf{r}_\mathbf{A}, \mathbf{r}'_\mathbf{A}$ of degree m and a random polynomial $\mathbf{u}_\mathbf{A}$ of degree $2m$. <ul style="list-style-type: none"> • Input $\mathbf{r}_\mathbf{A}, \mathbf{u}_\mathbf{A}$ into $\mathcal{F}_{\text{OPA}}^{(1)}$. • Input $\mathbf{p}_\mathbf{A}$ into $\mathcal{F}_{\text{OPA}}^{(2)}$ and obtain $\mathbf{s}_\mathbf{A} = \mathbf{p}_\mathbf{A}\mathbf{r}_\mathbf{B} + \mathbf{u}_\mathbf{B}$. • Set $\mathbf{s}'_\mathbf{A} = \mathbf{s}_\mathbf{A} - \mathbf{u}_\mathbf{A} + \mathbf{p}_\mathbf{A}\mathbf{r}'_\mathbf{A}$ and send it to Bob. 2. $P_\mathbf{B}$ (Input $S_\mathbf{B}$): Compute a polynomial $\mathbf{p}_\mathbf{B}$ of degree m such that $\mathbf{p}_\mathbf{B}(\gamma_j) = 0$ for all $\gamma_j \in S_\mathbf{B}$. Generate two random polynomials $\mathbf{r}_\mathbf{B}, \mathbf{r}'_\mathbf{B}$ of degree m and a random polynomial $\mathbf{u}_\mathbf{B}$ of degree $2m$. <ul style="list-style-type: none"> • Input $\mathbf{r}_\mathbf{B}, \mathbf{u}_\mathbf{B}$ into $\mathcal{F}_{\text{OPA}}^{(2)}$. • Input $\mathbf{p}_\mathbf{B}$ into $\mathcal{F}_{\text{OPA}}^{(1)}$ and obtain $\mathbf{s}_\mathbf{B} = \mathbf{p}_\mathbf{B}\mathbf{r}_\mathbf{A} + \mathbf{u}_\mathbf{A}$. • Upon receiving $\mathbf{s}'_\mathbf{A}$, compute $\mathbf{p}_\cap = \mathbf{s}'_\mathbf{A} + \mathbf{s}_\mathbf{B} + \mathbf{p}_\mathbf{B}\mathbf{r}'_\mathbf{B} - \mathbf{u}_\mathbf{B}$ and send it to Alice.
Output Verification
<ol style="list-style-type: none"> 3. Alice: Pick a random $x_\mathbf{A} \in \mathbb{F}$ and send it to Bob. 4. Bob: Set $\alpha_\mathbf{B} = \mathbf{p}_\mathbf{B}(x_\mathbf{A})$, $\beta_\mathbf{B} = \mathbf{r}_\mathbf{B}(x_\mathbf{A})$ and $\delta_\mathbf{B} = \mathbf{r}'_\mathbf{B}(x_\mathbf{A})$. Pick a random $x_\mathbf{B} \in \mathbb{F}$ and send $(x_\mathbf{B}, \alpha_\mathbf{B}, \beta_\mathbf{B}, \delta_\mathbf{B})$ to Alice. 5. Alice: Check if $\mathbf{p}_\mathbf{A}(x_\mathbf{A})(\beta_\mathbf{B} + \mathbf{r}'_\mathbf{A}(x_\mathbf{A})) + \alpha_\mathbf{B}(\mathbf{r}_\mathbf{A}(x_\mathbf{A}) + \delta_\mathbf{B}) = \mathbf{p}_\cap(x_\mathbf{A})$, otherwise abort. For each $\gamma_j \in S_\mathbf{A}$: If $\mathbf{p}_\cap(\gamma_j) = 0$, add γ_j to S_\cap. Send $\alpha_\mathbf{A} = \mathbf{p}_\mathbf{A}(x_\mathbf{B})$, $\beta_\mathbf{A} = \mathbf{r}_\mathbf{A}(x_\mathbf{B})$ and $\delta_\mathbf{A} = \mathbf{r}'_\mathbf{A}(x_\mathbf{B})$ to Bob. Output S_\cap. 6. Bob: Check if $\alpha_\mathbf{A}(\mathbf{r}_\mathbf{B}(x_\mathbf{B}) + \delta_\mathbf{A}) + \mathbf{p}_\mathbf{B}(x_\mathbf{B})(\beta_\mathbf{A} + \mathbf{r}'_\mathbf{B}(x_\mathbf{B})) = \mathbf{p}_\cap(x_\mathbf{B})$, otherwise abort. For each $\gamma_j \in S_\mathbf{B}$: If $\mathbf{p}_\cap(\gamma_j) = 0$, add γ_j to S_\cap. Output S_\cap.

Figure 9: Protocol $\Pi_{2\text{PSI}}$ UC-realises \mathcal{F}_{PSI} in the \mathcal{F}_{OPA} -hybrid model.

Theorem 1. *The protocol $\Pi_{2\text{PSI}}$ UC-realises \mathcal{F}_{PSI} in the \mathcal{F}_{OPA} -hybrid model with communication complexity $O(m\kappa)$.*

Proof. Let us argue that $\mathbf{p}_\cap = \mathbf{p}_\mathbf{A}(\mathbf{r}'_\mathbf{A} + \mathbf{r}_\mathbf{B}) + \mathbf{p}_\mathbf{B}(\mathbf{r}_\mathbf{A} + \mathbf{r}'_\mathbf{B})$ actually hides the inputs. The main observation here is that $\mathbf{r}'_\mathbf{P} + \mathbf{r}_\mathbf{P}$ is uniformly random as long as one party is honest. Since $\mathbf{p}_\mathbf{A} + \mathbf{p}_\mathbf{B}$

validly encodes the intersection (see Lemma 3), \mathbf{p}_\cap is uniformly random over the choice of the randomization polynomials $\mathbf{r}_A, \mathbf{r}'_A, \mathbf{r}_B$ and \mathbf{r}'_B , except for the roots denoting the intersection.

Corrupted Alice. We show the indistinguishability of the simulation of \mathcal{S}_A (cf. Figure 10). The simulator extracts Alice's inputs and then checks for any deviating behaviour. If such behaviour is detected, it aborts, even if the protocol would succeed. Proving indistinguishability of the simulation shows that the check in the protocol basically enforces semi-honest behaviour by Alice, up to input substitution.

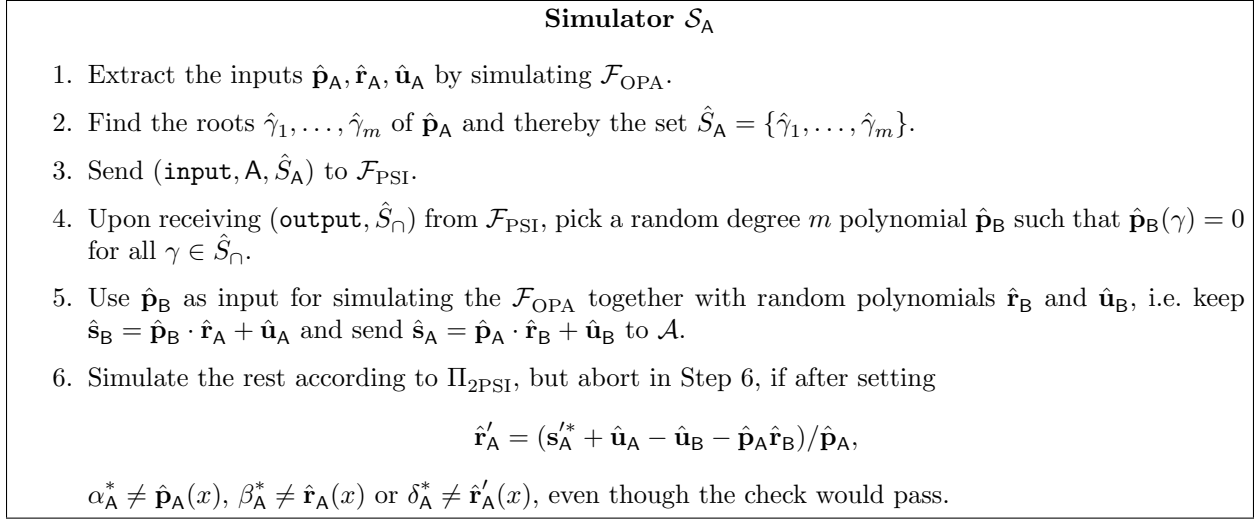


Figure 10: Simulator \mathcal{S}_A against a corrupted Alice.

Consider the following series of hybrid games.

Hybrid 0: $\text{Real}_{\Pi_{2\text{PSI}}}^{\mathcal{A}_A}$.

Hybrid 1: Identical to Hybrid 0, except that \mathcal{S}_1 simulates \mathcal{F}_{OPA} , learns all inputs and aborts if $\alpha_A^* \neq \hat{\mathbf{p}}_A(x)$ or $\beta_A^* \neq \hat{\mathbf{r}}_A(x)$, but the check is passed.

Let $\alpha_A^* = \alpha_A + e$ be \mathcal{A}_A 's check value. Then the check in Step 6 will fail with overwhelming probability. Let σ denote the outcome of the check. If \mathcal{A}_A behaves honestly, then

$$\sigma = \alpha_A^*(\mathbf{r}_B(x) + \delta_A^*) + \mathbf{p}_B(x)(\beta_A^* + \mathbf{r}'_B(x)) - \mathbf{p}_\cap(x) = 0.$$

Using $\alpha_A^* = \alpha_A + e$, however, we get

$$\sigma' = (\alpha_A + e)(\mathbf{r}_B(x) + \delta_A^*) + \mathbf{p}_B(x)(\beta_A^* + \mathbf{r}'_B(x)) - \mathbf{p}_\cap(x) = e \cdot (\mathbf{r}_B(x) + \delta_A^*) \neq \text{const.}$$

This means that the outcome of the check is uniformly random from \mathcal{A}_A 's view over the choice of \mathbf{r}_B (or \mathbf{p}_B for $\beta_A^* \neq \mathbf{r}_A(x)$). Therefore, the check will fail except with probability $2/|\mathbb{F}|$ and Hybrids 0 and 1 are statistically close.

Hybrid 2: Identical to Hybrid 1, except that \mathcal{S}_2 aborts according to Step 6 in Figure 10.

An environment distinguishing Hybrids 1 and 2 must manage to send \mathbf{s}'_A such that

$$\mathbf{s}'_A + \hat{\mathbf{u}}_A - \hat{\mathbf{u}}_B \neq \hat{\mathbf{p}}_A \cdot (\hat{\mathbf{r}}_B + \hat{\mathbf{r}}'_A)$$

while passing the check in Step 6 with non-negligible probability.

Let $\mathbf{f} = \mathbf{s}_A^* + \hat{\mathbf{u}}_A - \hat{\mathbf{u}}_B - \hat{\mathbf{p}}_A \cdot (\hat{\mathbf{r}}_B + \hat{\mathbf{r}}_A')$. We already know that $\mathbf{f}(x) = 0$, otherwise we have $\alpha_A^* = \alpha_A + \mathbf{f}(x) \neq \alpha_A$ (or an invalid β_A^*), and the check fails. But since x is uniformly random, the case that $\mathbf{f}(x) = 0$ happens only with probability $m/|\mathbb{F}|$, which is negligible. Therefore, Hybrid 1 and Hybrid 2 are statistically close.

Hybrid 3: Identical to Hybrid 2, except that \mathcal{S}_3 generates the inputs $\hat{\mathbf{s}}_A, \hat{\mathbf{s}}_B$ according to Step 5 in Figure 10 and adjusts the output. This corresponds to $\text{Ideal}_{\mathcal{F}_{\text{PSI}}}^{\mathcal{S}_A}$.

The previous hybrids established that the inputs $\hat{\mathbf{p}}_A, \hat{\mathbf{r}}_A$ are extracted correctly. Therefore, by definition, $\hat{\mathcal{S}}_A = \mathcal{M}_0(\hat{\mathbf{p}}_A)$. It remains to argue that the simulated outputs are indistinguishable. First, note that the received intersection $\hat{\mathcal{S}}_\cap = \mathcal{M}_0(\hat{\mathbf{p}}_B)$ defines $\hat{\mathbf{p}}_B$. From Lemma 6 it follows that $\mathcal{M}_0(\mathbf{p}_\cap) = \mathcal{M}_0(\hat{\mathbf{p}}_A) \cap \mathcal{M}_0(\hat{\mathbf{p}}_B) = \hat{\mathcal{S}}_\cap$ w.r.t. $\mathcal{M}_0(\hat{\mathbf{p}}_B)$, even for a maliciously chosen $\hat{\mathbf{r}}_A$, i.e. the \mathcal{A}_A cannot increase the intersection even by a single element except with negligible probability.

Further, note that $\hat{\mathbf{s}}_A = \hat{\mathbf{p}}_A \cdot \hat{\mathbf{r}}_B + \hat{\mathbf{u}}_B$ is uniformly distributed over the choice of $\hat{\mathbf{u}}_B$, and $\hat{\mathbf{p}}_\cap$ is uniform over the choice of $\hat{\mathbf{r}}_B, \hat{\mathbf{r}}_B'$.

Finally, since $\hat{\mathbf{r}}_B, \hat{\mathbf{r}}_B'$ are uniformly random and the degree of $\hat{\mathbf{p}}_B$ is m , i.e. $\max_i |\mathcal{S}_i| + 1$, the values $\hat{\alpha}_B, \hat{\beta}_B$ and $\hat{\delta}_B$ are uniformly distributed as well. In conclusion, the Hybrids 2 and 3 are statistically close.

As a result we get that for all environments \mathcal{Z} ,

$$\text{Real}_{\Pi_{2\text{PSI}}}^{\mathcal{A}_A}(\mathcal{Z}) \approx_s \text{Ideal}_{\mathcal{F}_{\text{PSI}}}^{\mathcal{S}_A}(\mathcal{Z}).$$

Corrupted Bob. The simulator against a corrupted Bob in Figure 11 (and therefore the proof) is essentially the same as the one against a corrupted Alice, except for a different way to check his inputs.

Simulator \mathcal{S}_B

1. Extract the inputs $\hat{\mathbf{p}}_B, \hat{\mathbf{r}}_B, \hat{\mathbf{u}}_B$ by simulating \mathcal{F}_{OPA} .
2. Find the roots $\hat{\gamma}_1, \dots, \hat{\gamma}_m$ of $\hat{\mathbf{p}}_B$ and thereby the set $\hat{\mathcal{S}}_B = \{\hat{\gamma}_1, \dots, \hat{\gamma}_m\}$.
3. Send $(\text{input}, B, \hat{\mathcal{S}}_B)$ to \mathcal{F}_{PSI} .
4. Upon receiving $(\text{output}, \hat{\mathcal{S}}_\cap)$ from \mathcal{F}_{PSI} , pick a random degree m polynomials $\hat{\mathbf{p}}_A$ such that $\hat{\mathbf{p}}_A(\gamma) = 0$ for all $\gamma \in \hat{\mathcal{S}}_\cap$.
5. Use $\hat{\mathbf{p}}_A$ as input for simulating the \mathcal{F}_{OPA} together with random polynomials $\hat{\mathbf{r}}_A$ and $\hat{\mathbf{u}}_A$, i.e. send $\hat{\mathbf{s}}_B = \hat{\mathbf{p}}_B \cdot \hat{\mathbf{r}}_A + \hat{\mathbf{u}}_A$ to \mathcal{A}_B and keep $\hat{\mathbf{s}}_A = \hat{\mathbf{p}}_A \cdot \hat{\mathbf{r}}_B + \hat{\mathbf{u}}_B$.
6. Simulate the rest according to $\Pi_{2\text{PSI}}$, but abort in Step 5, if after setting

$$\hat{\mathbf{r}}_A' = \mathbf{p}_\cap^* - (\hat{\mathbf{p}}_A(\hat{\mathbf{r}}_B + \hat{\mathbf{r}}_A') + \hat{\mathbf{p}}_B(\hat{\mathbf{r}}_B' + \hat{\mathbf{r}}_A)),$$
 the extracted $\alpha_B^* \neq \hat{\mathbf{p}}_B(x)$, $\beta_B^* \neq \hat{\mathbf{r}}_B(x)$ or $\delta_B^* \neq \hat{\mathbf{r}}_B'(x)$, even if the check passes otherwise.

Figure 11: Simulator \mathcal{S}_B against a corrupted Bob.

Consider the following series of hybrid games.

Hybrid 0: $\text{Real}_{\Pi_{2\text{PSI}}}^{\mathcal{A}_B}$.

Hybrid 1: Identical to Hybrid 0, except that \mathcal{S}_1 simulates \mathcal{F}_{OPA} , learns all inputs and aborts if $\alpha_A^* \neq \hat{\mathbf{p}}_A(x)$ or $\beta_A^* \neq \hat{\mathbf{r}}_A(x)$, but the check is passed.

This step is identical to the case of a corrupted Alice. Let $\alpha_B^* = \alpha_B + e$ be \mathcal{A}_B 's check value. Then the check in Step 5 will fail with overwhelming probability. Let σ denote the outcome of the check. If \mathcal{A}_B behaves honestly, then

$$\sigma = \alpha_B^*(\mathbf{r}_A(x) + \delta_B^*) + \mathbf{p}_A(x)(\beta_B^* + \mathbf{r}'_A(x)) - \mathbf{p}_\cap(x).$$

Using $\alpha_B^* = \alpha_B + e$, however, we get

$$\sigma' = (\alpha_B + e)(\mathbf{r}_A(x) + \delta_B^*) + \mathbf{p}_A(x)(\beta_B^* + \mathbf{r}'_A(x)) - \mathbf{p}_\cap(x) = e \cdot (\mathbf{r}_A(x) + \delta_B^*) \neq \text{const.}$$

This means that the outcome of the check is uniformly random from \mathcal{A}_B 's view over the choice of \mathbf{r}_A (or \mathbf{p}_A for $\beta_B^* \neq \mathbf{r}_B(x)$). Therefore, the check will fail except with probability $2/|\mathbb{F}|$ and Hybrids 0 and 1 are statistically close.

Hybrid 2: Identical to Hybrid 1, except that \mathcal{S}_2 aborts according to Step 6 in Figure 11.

An environment distinguishing Hybrids 1 and 2 must manage to send \mathbf{p}_\cap^* such that

$$\mathbf{p}_\cap^* \neq \hat{\mathbf{p}}_A(\hat{\mathbf{r}}_B + \hat{\mathbf{r}}'_A) + \hat{\mathbf{p}}_B(\hat{\mathbf{r}}'_B + \hat{\mathbf{r}}_A),$$

while passing the check in Step 5 with non-negligible probability.

Let $\mathbf{f} = \mathbf{p}_\cap^* - (\hat{\mathbf{p}}_A(\hat{\mathbf{r}}_B + \hat{\mathbf{r}}'_A) + \hat{\mathbf{p}}_B(\hat{\mathbf{r}}'_B + \hat{\mathbf{r}}_A))$. We already know that $\mathbf{f}(x) = 0$, otherwise we have $\alpha_B^* = \alpha_B + \mathbf{f}(x) \neq \alpha_B$ (or an invalid β_B^*), and the check fails. But since x is uniformly random, the case that $\mathbf{f}(x) = 0$ happens only with probability $m/|\mathbb{F}|$, which is negligible. Therefore, Hybrid 1 and Hybrid 2 are statistically close.

Hybrid 3: Identical to Hybrid 2, except that \mathcal{S}_3 generates the inputs $\hat{\mathbf{s}}_B, \hat{\mathbf{s}}_A$ according to Step 5 in Figure 11 and adjusts the output. This corresponds to $\text{Ideal}_{\mathcal{F}_{\text{PSI}}}^{\mathcal{S}_B}$.

The previous hybrids established that the inputs $\hat{\mathbf{p}}_B, \hat{\mathbf{r}}_B$ are extracted correctly. Therefore, by definition, $\hat{S}_B = \mathcal{M}_0(\hat{\mathbf{p}}_B)$. It remains to argue that the simulated outputs are indistinguishable. First, note that the received intersection $\hat{S}_\cap = \mathcal{M}_0(\hat{\mathbf{p}}_A)$ defines $\hat{\mathbf{p}}_A$. From Lemma 6 it follows that $\mathcal{M}_0(\mathbf{p}_\cap) = \mathcal{M}_0(\hat{\mathbf{p}}_B) \cap \mathcal{M}_0(\hat{\mathbf{p}}_A) = \hat{S}_\cap$ w.r.t. $\mathcal{M}_0(\hat{\mathbf{p}}_A)$, even for a maliciously chosen $\hat{\mathbf{r}}_B$, i.e. \mathcal{A}_B cannot increase the intersection even by a single element except with negligible probability.

Further, note that $\hat{\mathbf{s}}_B = \hat{\mathbf{p}}_B \cdot \hat{\mathbf{r}}_A + \hat{\mathbf{u}}_A$ is uniformly distributed over the choice of $\hat{\mathbf{u}}_A$, and $\hat{\mathbf{p}}_\cap$ is uniform over the choice of $\hat{\mathbf{r}}_A$.

Finally, since $\hat{\mathbf{r}}_A, \hat{\mathbf{r}}'_A$ are uniformly random and the degree of $\hat{\mathbf{p}}_A$ is m , i.e. $\max_i |\mathcal{S}_i| + 1$, the values $\hat{\alpha}_A, \hat{\beta}_A$ and $\hat{\delta}_A$ are uniformly distributed as well. In conclusion, the Hybrids 2 and 3 are statistically close.

As a result we get that for all environments \mathcal{Z} ,

$$\text{Real}_{\Pi_{2\text{PSI}}}^{\mathcal{A}_B}(\mathcal{Z}) \approx_s \text{Ideal}_{\mathcal{F}_{\text{PSI}}}^{\mathcal{S}_B}(\mathcal{Z}).$$

Efficiency. The protocol makes two calls to OPA, which in turn is based on OLE. Overall, $2m$ calls to OLE are necessary in OPA. Given the recent constant overhead OLE of Ghosh et al. [GNN17], the communication complexity of $\Pi_{2\text{PSI}}$ lies in $O(m)$.

On the computational side, the parties have to compute interpolations of polynomials of degree m , which brings the computational complexity to $O(m \log m)$ using NTT. Note that this cost includes the computational cost of the OLE instantiation from [GNN17]. This concludes the proof. \square

6 Maliciously Secure Multi-party PSI

6.1 Ideal Functionality

The ideal functionality for multi-party private set intersection $\mathcal{F}_{\text{MPSI}}^*$ simply takes the inputs from all parties and computes the intersection of these inputs. Our functionality $\mathcal{F}_{\text{MPSI}}^*$ in Figure 12 additionally allows an adversary to learn the intersection and then possibly update the result to be only a subset of the original result.




Figure 12: Ideal functionality $\mathcal{F}_{\text{MPSI}}^*$ for multi-party PSI.

Let us briefly elaborate on why we chose to use this modified functionality. In the UC setting, in order to extract the inputs of all malicious parties, any honest party has to communicate with all malicious parties. In particular, since the simulator has to extract the complete input, this requires at least $O(nm)$ communication per party for the classical MPSI functionality. In turn, for the complete protocol, this means that the communication complexity lies in $O(n^2m)$.

Instead, we want to take an approach similar to the recent work of Hazay et al. [HV17], i.e. we have one central party, and some of the work is delegated to this party. This removes the need for the other parties to extensively communicate with each other and potentially allows communication complexity $O(mn)$, which is asymptotically optimal in any setting. However, if we assume that the central party and at least one additional party are corrupted, the honest party does not (extensively) interact with this additional party and does not learn its inputs; it can only learn the input of the central party. If the input set of the other malicious party is the same as the one of the central party, the output remains the same. If this input is different, however, the actual intersection might be smaller. One might argue that this case simply corresponds to input substitution by the malicious party, but for any type of UC simulation this poses a problem, since the output of the honest party in the protocol might be different from the intersection in the ideal world. Thus, $\mathcal{F}_{\text{MPSI}}^*$ allows a malicious party to modify the output. Crucially, the updated intersection can only be smaller

and may not be changed arbitrarily by the adversary. We believe that this weaker multiparty PSI functionality is sufficient for most scenarios.

6.2 Multi-party PSI from OLE

Our multi-party PSI protocol uses the same techniques that we previously employed to achieve two-party PSI. This is similar in spirit to the approach taken in [HV17], who employ techniques from the two-party PSI of [FNP04] and apply them in the multi-party setting. We also adopt the idea of a designated central party that performs a two-party PSI with the remaining parties, because this allows to delegate most of the computation to this party and saves communication. Apart from that, our techniques differ completely from [HV17]. Abstractly, they run the two-party PSI with each party and then use threshold encryption and zero-knowledge proofs to ensure the correctness of the computation. These tools inflict a significant communication and computation penalty.

In our protocol (cf. Figure 13) we run our two-party PSI between the central party and every other party, but we ensure privacy of the aggregation not via threshold encryption and zero-knowledge proofs, but instead by a simple masking of the intermediate values and a polynomial check. This masking is created in a setup phase, where every pair of parties exchanges a random seed that is used to create two random blinding polynomials which cancel out when added.

Once the central party receives all shares of the computation, it simply adds these shares, thereby removing the random masks. The central party broadcasts the result to all parties. Then, all parties engage in a multi-party coin-toss and obtain a random value x . Since all operations in the protocol are linear operations on polynomials, the parties evaluate their input polynomials on x and broadcast the result. This allows every party to locally verify the relation and as a consequence also the result. Here we have to ensure that a rushing adversary cannot cheat by waiting for all answers before providing its own answer. We solve this issue by simply committing to the values first, and the unveiling them in the next step. This leads to malleability problems, i.e. we have to use non-malleable commitments³.

Let us briefly give an intuition on why this protocol computes the intersection of all parties. The intersection polynomial $\mathbf{p}_\cap = \sum_{i=1}^{n-1} (\mathbf{p}_0 + \mathbf{p}_i) \cdot (\mathbf{r}_i + \mathbf{r}_0^i)$ is the sum of all two-party intersection polynomials. Every such intermediate polynomial contains exactly the intersection between the parties P_i and P_0 in its roots (plus some additional, but random, roots). By simply adding all of these polynomials, the common roots of the intermediate polynomials are preserved, while the other roots are blinded by random values. The probability that two of these random blindings cancel out is negligible in the field size. Therefore, from the view of each party, the common roots of \mathbf{p}_\cap and \mathbf{p}_i represent the intersection with all other parties.

Theorem 2. *The protocol Π_{MPSI} computationally UC-realises $\mathcal{F}_{\text{MPSI}}^*$ in the \mathcal{F}_{OPA} -hybrid model with communication complexity in $O((n^2 + nm)\kappa)$.*

Proof. We have to distinguish between the case where the central party is malicious and the case where it is honest. We show UC-security of Π_{MPSI} by defining a simulator \mathcal{S} for each case which produces an indistinguishable simulation of the protocol to any environment \mathcal{Z} trying to distinguish

³In order to achieve our claimed efficiency we actually use UC commitments, but non-malleable commitments are sufficient for the security of the protocol.

$$\Pi_{\text{MPSI}}$$

Let $m = \max_i \{|\mathcal{S}_i|\} + 1$ and NMCOM be a bounded-concurrent non-malleable commitment scheme against synchronized adversaries. $\mathcal{F}_{\text{OPA}}^{(i,j)}$ denotes the j th instance for parties P_0 and P_i .

Setup

1. All parties P_i and P_j for $i, j \in \{1, \dots, n-1\}$ exchange a random polynomial as follows. For all $j \neq i$, if $\mathbf{v}_{ij} = \perp$, P_i picks seed_{ij} uniformly at random and sets $\mathbf{v}_{ij} = \text{PRG}_{2m}(\text{seed}_{ij})$. It sends seed_{ij} to P_j , who sets $\mathbf{v}_{ji} = -\text{PRG}_{2m}(\text{seed}_{ij})$.

Share Computation

2. P_0 (Input S_0): Compute a polynomial \mathbf{p}_0 of degree m s.t. $\mathbf{p}_0(\gamma_j) = 0$ for all $\gamma_j \in S_0$. Generate n random polynomials $\mathbf{r}_0^i \in \mathbb{F}[X], i \in \{1, \dots, n-1\}$ and $\mathbf{r}'_0 \in \mathbb{F}[X]$ of degree m each and $n-1$ random polynomials $\mathbf{u}_0^i \in \mathbb{F}[X], i \in \{1, \dots, n-1\}$ of degree $2m$. For $i \in [n-1]$
 - Input $\mathbf{r}_0^i, \mathbf{u}_0^i$ into $\mathcal{F}_{\text{OPA}}^{(i,1)}$ for each $i \in \{1, \dots, n-1\}$.
 - Input \mathbf{p}_0 into $\mathcal{F}_{\text{OPA}}^{(i,2)}$ and obtain $\mathbf{s}_0^i = \mathbf{p}_0 \cdot \mathbf{r}_i + \mathbf{u}_i$.
3. P_i (Input S_i): Compute a polynomial \mathbf{p}_i of degree m s.t. $\mathbf{p}_i(\gamma_j) = 0$ for all $\gamma_j \in S_i$. Additionally, pick $\mathbf{r}_i, \mathbf{r}'_i \in \mathbb{F}[X]$ uniformly of degree m and $\mathbf{u}_i \in \mathbb{F}[X]$ uniformly of degree $2m$.
 - Input \mathbf{p}_i into $\mathcal{F}_{\text{OPA}}^{(i,1)}$ and obtain $\mathbf{s}_i = \mathbf{p}_i \cdot \mathbf{r}_0^i + \mathbf{u}_0^i$.
 - Input $\mathbf{r}_i, \mathbf{u}_i$ into $\mathcal{F}_{\text{OPA}}^{(i,2)}$.
 - Set $\mathbf{s}'_i = \mathbf{s}_i - \mathbf{u}_i + \mathbf{p}_i \mathbf{r}'_i + \sum_{i \neq j} \mathbf{v}_{ij}$ and send it to P_0 .

Output Aggregation and Verification

4. P_0 : Compute $\mathbf{p}_\cap = \sum_{i=1}^{n-1} (\mathbf{s}'_i + \mathbf{s}_0^i - \mathbf{u}_0^i + \mathbf{p}_0 \mathbf{r}'_0)$ and broadcast \mathbf{p}_\cap .
5. *All parties*:
 - Run a multiparty coin-toss protocol Π_{CT} to obtain a random $x \in \mathbb{F}$.
 - Evaluate $\alpha_i = \mathbf{p}_i(x)$, $\beta_i = \mathbf{r}_i(x)$, $\delta_i = \mathbf{r}'_i(x)$ and compute $(\text{com}_i, \text{unv}_i) = \text{NMCOM.Commit}(\alpha_i, \beta_i, \delta_i)$. Broadcast com_i .
 - Once all commitments are received, broadcast unv_i and $(\alpha_i, \beta_i, \delta_i)$. Abort if $\sum \alpha_0 \cdot (\beta_i + \delta_0) + \alpha_i \cdot (\beta_0 + \delta_i) \neq \mathbf{p}_\cap(x)$ or $\text{NMCOM.Open}(\text{com}_i, \text{unv}_i, (\alpha_i, \beta_i, \delta_i)) \neq 1$.
 - For each $\gamma_j \in S_i$: if $\mathbf{p}_\cap(\gamma_j) = 0$, add γ_j to S_\cap . Output S_\cap .

Figure 13: Protocol Π_{MPSI} UC-realises $\mathcal{F}_{\text{MPSI}}^*$ in the \mathcal{F}_{OPA} -hybrid model.

the ideal world from the real world. The approach of the simulation is straightforward: the simulator extracts the input polynomials into \mathcal{F}_{OPA} and thus obtains an intersection of the adversary's inputs.

In the case of an honest central party, all parties communicate with this party, i.e. the simulator can extract all inputs of all malicious parties. In the case where P_0 is malicious, however, the simulator can at most learn the central party's input at the beginning. He inputs this result into the ideal functionality and uses the intermediate result for the simulation. The malicious central party can later "simulate" the other malicious parties and thereby possibly change the intersection

for the honest parties. We show that \mathcal{A} can only reduce the intersection unless it already knows $x \in S_j$ for at least one $j \in \mathbf{H}$, i.e. we assume that \mathcal{A} cannot predict a single element of the set of an honest party except with negligible probability. This reduced intersection can be passed by the simulator to the ideal functionality.

P_0 is malicious: Consider the simulator in Figure 14.

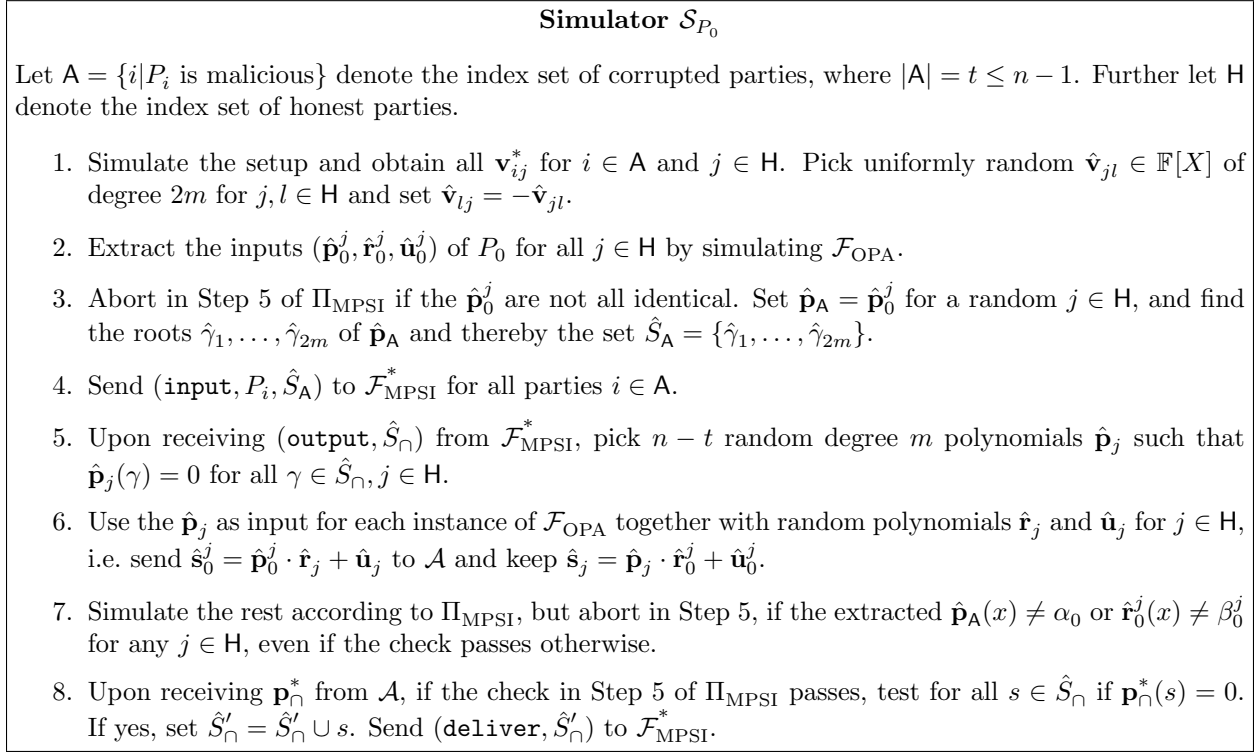


Figure 14: Simulator \mathcal{S}_{P_0} for $P_0 \in \mathbf{A}$.

We show the indistinguishability of the simulation and the real protocol through the following hybrid games. In the following, let \mathcal{A} denote the dummy adversary controlled by \mathcal{Z} .

Hybrid 0: $\text{Real}_{\Pi_{\text{MPSI}}}^{\mathbf{A}}$.

Hybrid 1: Identical to Hybrid 0, except that \mathcal{S}_1 simulates \mathcal{F}_{OPA} and learns all inputs.

Hybrid 2: Identical to Hybrid 1, except that \mathcal{S}_2 aborts according to Step 7 in Figure 14.

Hybrid 3: Identical to Hybrid 2, except that \mathcal{S}_3 aborts if the extracted $\hat{\mathbf{p}}_0$ are not identical, but the check is passed.

Hybrid 4: Identical to Hybrid 3, except that \mathcal{S}_4 replaces the \mathbf{v}_{jl} between honest parties j, l by uniformly random polynomials.

Hybrid 5: Identical to Hybrid 4, except that \mathcal{S}_5 generates the inputs $\hat{\mathbf{s}}_0^j, \hat{\mathbf{s}}_j$ according to Step 6 in Figure 14 and adjusts the output. This corresponds to $\text{Ideal}_{\mathcal{F}_{\text{MPSI}}^*}^{\mathcal{S}_{P_0}}$.

Hybrids 0 and 1 are trivially indistinguishable. We show that Hybrid 1 and Hybrid 2 are computationally indistinguishable in Lemma 9.1. This step ensures that the correct \hat{p}_0 was extracted, and that all the intermediate values of the honest parties are added up. Hybrids 2 and 3 are indistinguishable due to the security of the coin-toss. This is formalized in Lemma 9.2. As an intermediate step to complete the full simulation, we replace all pseudorandom polynomials \mathbf{v}_{jl} between honest parties j, l by uniformly random ones. Computational indistinguishability of Hybrid 3 and Hybrid 4 follows from a straightforward reduction to the pseudorandomness of PRG. We establish the statistical indistinguishability of Hybrids 4 and 5 in Lemma 9.3. As a result we get that for all PPT environments \mathcal{Z} ,

$$\text{Real}_{\Pi_{\text{MPSI}}}^{\mathcal{A}}(\mathcal{Z}) \approx_c \text{Ideal}_{\mathcal{F}_{\text{MPSI}}}^{\mathcal{S}_{P_0}}(\mathcal{Z}).$$

Lemma 9.1. *Assume that NMCOM is a bounded-concurrent non-malleable commitment scheme against synchronizing adversaries. Then Hybrid 1 and Hybrid 2 are computationally indistinguishable.*

Proof. The only difference between Hybrid 1 and Hybrid 2 lies in the fact that \mathcal{S}_2 aborts if the extracted $\hat{\mathbf{p}}_{\mathbf{A}}$ evaluated on x does not match the check value α_0 , but the check is still passed. Therefore, in order for \mathcal{Z} to distinguish both hybrids, it has to be able to produce a value $\alpha_0^* \neq \hat{\mathbf{p}}_{\mathbf{A}}(x)$ and pass the check with non-negligible probability ϵ . W.l.o.g. it is sufficient that α_0^* is incorrect for only one $\hat{\mathbf{p}}_0$. We show that such a \mathcal{Z} breaks the non-malleability property of NMCOM.

Let σ denote the outcome of the check. If \mathcal{A} is honest, i.e. $\alpha_0 = \hat{\mathbf{p}}_0(x)$ and $\beta_0^i = \hat{\mathbf{r}}_0^i(x)$, then

$$\sigma = \sum_{i=0}^n (\alpha_0(\beta_i + \delta_0) + \alpha_i(\beta_0^i + \delta_i)) - \mathbf{p}_{\cap}(x) = 0, \quad (1)$$

where

$$\mathbf{p}_{\cap} = \sum_{i \in \mathbf{A}} (\mathbf{s}_i + \mathbf{s}_0^i) + \sum_{j \in \mathbf{H}} (\mathbf{s}_j + \mathbf{s}_0^j).$$

We first observe that $\sum_{j \in \mathbf{H}} (\mathbf{s}_j + \mathbf{s}_0^j) = \sum_{j \in \mathbf{H}} \hat{\mathbf{p}}_j(\hat{\mathbf{r}}_0^j + \hat{\mathbf{r}}_j^j) + \hat{\mathbf{p}}_0(\hat{\mathbf{r}}_0^j + \hat{\mathbf{r}}_j^j)$ is uniform over the choice of the $\hat{\mathbf{r}}_j, \hat{\mathbf{r}}_j'$. Therefore, if \mathcal{A} uses \mathbf{p}_{\cap}^* without adding $\sum_{j \in \mathbf{H}} (\mathbf{s}_j + \mathbf{s}_0^j)$, the check will fail with overwhelming probability.

Since \mathcal{A} controls the inputs of the malicious parties $i \in \mathbf{A}$, in order to pass the check it is sufficient for \mathcal{A} to satisfy the following simplification of Equation (1).

$$\sigma' = \sum_{j \in \mathbf{H}} (\alpha_0(\beta_j + \delta_0) + \alpha_j(\beta_0^j + \delta_j)) - \sum_{j \in \mathbf{H}} (\mathbf{s}_j(x) + \mathbf{s}_0^j(x)) = \text{const}$$

Here const is a fixed constant known to \mathcal{A} (0 if \mathcal{A} is honest) determined by setting the inputs α_i, β_i for $i \in \mathbf{A}$ accordingly. But if $\alpha_0^* \neq \hat{\mathbf{p}}_0(x)$, i.e. $\alpha_0^* = \alpha_0 + e$, then we get that

$$\begin{aligned} \sigma' &= \sum_{j \in \mathbf{H}} ((\alpha_0 + e)(\beta_j + \delta_0) + \alpha_j(\beta_0^j + \delta_j)) - \sum_{j \in \mathbf{H}} (\mathbf{s}_j(x) + \mathbf{s}_0^j(x)) \\ &= \sum_{j \in \mathbf{H}} (\alpha_0(\beta_j + \delta_0) + \alpha_j(\beta_0^j + \delta_j)) - \sum_{j \in \mathbf{H}} (\mathbf{s}_j(x) + \mathbf{s}_0^j(x)) + e \sum_{j \in \mathbf{H}} (\beta_j + \delta_0) \\ &= e \sum_{j \in \mathbf{H}} (\beta_j + \delta_0) \neq \text{const} \end{aligned}$$

Similarly for $\beta_0^j \neq \hat{\mathbf{r}}_0^j(x)$ for any $j \in \mathbf{H}$. Thus, except for the case of $\alpha_0^* = \alpha_0 + e / \sum_{j \in \mathbf{H}} \beta_j$, the check will fail for $\alpha_0^* \neq \hat{\mathbf{p}}_0(x)$. But since we assumed that \mathcal{A} passes the check with non-negligible probability, and NMCOM is statistically binding, \mathcal{A} has to produce a valid commitment to $\tilde{\alpha}_0 = \alpha_0 + e / \sum_{j \in \mathbf{H}} (\beta_j + \delta_0)$ with the same probability.

Note, that \mathcal{A} interacts in both the left and right session of NMCOM with the same party (actually all parties simultaneously, since everything is broadcast). But this means that \mathcal{A} cannot let the left session finish before starting the right session, i.e. \mathcal{A} is a synchronizing adversary against NMCOM. Concretely, in the left session, \mathcal{S}_2 commits to $(\hat{\mathbf{p}}_j(x), \hat{\mathbf{r}}_j(x), \hat{\mathbf{r}}_j'(x)) = (\alpha_j, \beta_j, \delta_j)$ for $j \in \mathbf{H}$, while \mathcal{A} commits in the right session to $(\alpha_0, \{\beta_0^i\}_{i \in [n]}, \delta_0)$ and $(\alpha_i, \beta_i, \delta_i)$ for $i \in \mathbf{A}$ to \mathcal{S}_2 . Further, the number of sessions that \mathcal{A} can start is bounded in advance at $n - 1$, i.e. it is sufficient to consider bounded-concurrency.

Consider the two views

$$\text{Real} = \{\hat{\mathbf{s}}_j, \{\text{com}_j\}\}_{j \in \mathbf{H}}, \quad \text{Rand} = \{\hat{\mathbf{s}}_j, \{\widehat{\text{com}}_j\}\}_{j \in \mathbf{H}},$$

where $\text{com}_j \leftarrow \text{NMCOM.Commit}(\alpha_j, \beta_j)$ and $\widehat{\text{com}}_j \leftarrow \text{NMCOM.Commit}(0)$. Real corresponds to a real protocol view of \mathcal{A} before committing itself⁴.

Obviously, $\text{Real} \approx_c \text{Rand}$ if NMCOM is non-malleable. However, we will argue that \mathcal{A} cannot output a valid commitment on $\tilde{\alpha}_0$ except with negligible probability, i.e.

$$\Pr[(\text{com}_0^*, \text{unv}_0^*, (\tilde{\alpha}_0, \{\tilde{\beta}_0^i\}_{i \in [n]}, \tilde{\delta}_0) \leftarrow \mathcal{A}(\text{Rand}) \wedge \text{valid}] \leq \text{negl}(\kappa),$$

where valid is the event that $\text{NMCOM.Open}(\text{com}_0^*, \text{unv}_0^*, (\tilde{\alpha}_0, \{\tilde{\beta}_0^i\}_{i \in [n]}, \tilde{\delta}_0) = 1$. We first observe that $\hat{\mathbf{p}}_j$ and $\hat{\mathbf{r}}_j$ for $j \in \mathbf{H}$ cannot be obtained by \mathcal{A} via $\hat{\mathbf{s}}_j = \hat{\mathbf{p}}_j \cdot \hat{\mathbf{r}}_0^j - \hat{\mathbf{u}}_j$. The polynomial $\hat{\mathbf{s}}_j$ itself is uniformly random over the choice of $\hat{\mathbf{u}}_j$, and the only equation that \mathcal{A} has is $\hat{\mathbf{p}}_\cap = \sum_{i \in \mathbf{A}} (\mathbf{s}_i + \mathbf{s}_0^i) + \sum_{j \in \mathbf{H}} (\mathbf{s}_j + \mathbf{s}_0^j) = \sum_{i \in \mathbf{A}} (\hat{\mathbf{p}}_0 \cdot (\hat{\mathbf{r}}_i + \hat{\mathbf{r}}_0') + \hat{\mathbf{p}}_i \cdot (\hat{\mathbf{r}}_0^i + \hat{\mathbf{r}}_i')) + \sum_{j \in \mathbf{H}} (\hat{\mathbf{p}}_0 \cdot (\hat{\mathbf{r}}_j + \hat{\mathbf{r}}_0') + \hat{\mathbf{p}}_j \cdot (\hat{\mathbf{r}}_0^j + \hat{\mathbf{r}}_j'))$. Note, that the honest $\hat{\mathbf{r}}_j, \hat{\mathbf{r}}_j'$ have degree d and therefore hide $\hat{\mathbf{p}}_j$. Further, the commitments com_j contain the value 0 and are therefore independent of $\hat{\mathbf{p}}_j$ and $\hat{\mathbf{r}}_j$. Thus, the probability that \mathcal{A} obtains a commitment on $\tilde{\alpha}_0$ is negligible.

But since $\text{Real} \approx_c \text{Rand}$, we also get that

$$\Pr[(\text{com}_0^*, \text{unv}_0^*, (\tilde{\alpha}_0, \{\tilde{\beta}_0^i\}_{i \in [n]}, \tilde{\delta}_0) \leftarrow \mathcal{A}(\text{Real}) \wedge \text{valid}] \leq \text{negl}(\kappa),$$

which contradicts our assumption that \mathcal{A} produces the commitment with non-negligible probability ϵ .

In conclusion, Hybrid 1 and Hybrid 2 are computationally indistinguishable. \square

Lemma 9.2. *Assume that Π_{CT} provides a uniformly random x with computational security. Then Hybrid 2 and Hybrid 3 are computationally indistinguishable.*

Proof. Assume that there exists an environment \mathcal{Z} that distinguishes Hybrids 2 and 3 with non-negligible probability ϵ . In order to distinguish Hybrid 2 and Hybrid 3 \mathcal{Z} has to provide two distinct

⁴For ease of notation, here we assume that the commitments are completely sent before \mathcal{A} commits himself. The very same argument also holds if \mathcal{A} only received synchronized messages of com_j and has to start committing concurrently.

polynomials for a malicious P_0 and still pass the check in the protocol. Then we can construct from \mathcal{Z} an adversary \mathcal{B} that predicts the outcome of Π_{CT} with non-negligible probability.

Let \mathcal{A} input w.l.o.g. two polynomials $\hat{\mathbf{p}}_0^1 \neq \hat{\mathbf{p}}_0^2$. The check with the random challenge x allows \mathcal{A} to send only one value α_0^* , but from Lemma 9.1 we know that it has to hold that $\alpha_0^* = \hat{\mathbf{p}}_0^1(x) = \hat{\mathbf{p}}_0^2(x)$, or the check will fail. First note that two polynomials of degree m agree in a random point x over \mathbb{F} only with probability $m/|\mathbb{F}|$, which is negligible in our case.

Our adversary \mathcal{B} proceeds as follows. It simulates the protocol for \mathcal{Z} according to \mathcal{S}_1 up to the point where \mathcal{S}_1 learns the polynomials $\hat{\mathbf{p}}_0^1 \neq \hat{\mathbf{p}}_0^2$. \mathcal{B} sets $\mathbf{f} = \hat{\mathbf{p}}_0^1 - \hat{\mathbf{p}}_0^2$ and computes the roots $\gamma_1, \dots, \gamma_m$ of \mathbf{f} . One of these roots has to be the random point x , otherwise $\hat{\mathbf{p}}_0^1(x) - \hat{\mathbf{p}}_0^2(x) \neq 0$ and the check in Π_{MPSI} fails (since there is only one α_0^*). \mathcal{B} picks a random index $l \in [m]$ and predicts the output of the coin-flip as γ_l . Thus, \mathcal{B} predicts the outcome of the coin-toss correctly with probability ϵ/m , which is non-negligible. This contradicts the security of Π_{CT} .

This establishes the indistinguishability of Hybrid 2 and Hybrid 3. \square

Lemma 9.3. *Hybrid 4 and Hybrid 5 are statistically close.*

Proof. A malicious environment \mathcal{Z} can distinguish Hybrid 4 and Hybrid 5 if (a) the extracted inputs are incorrect or if (b) the simulated messages can be distinguished from real ones.

Concerning (a), if the inputs were not correctly extracted, \mathcal{Z} would receive different outputs in the two hybrids. We already established that the extracted polynomial $\hat{\mathbf{p}}_0$ is correct. Similarly, the extracted $\hat{\mathbf{r}}_0^j$ are also correct. By implication this also ensures that the intermediate intersection is computed correctly.

We argue that the correction of the intersection is also correct, i.e. the set \hat{S}'_\cap is computed correctly and in particular it holds that $(\mathcal{M}_0(\mathbf{p}_\cap^*) \cap \mathcal{M}_0(\hat{\mathbf{p}}_j)) \subseteq \hat{S}'_\cap$. First of all, we have to show that the intermediate intersection polynomial $\hat{\mathbf{p}}_{\text{int}}$ actually provides the intersection for all parties. For all P_j it holds with overwhelming probability:

$$\begin{aligned}
\mathcal{M}_0(\hat{\mathbf{p}}_j) \cap \mathcal{M}_0(\hat{\mathbf{p}}_{\text{int}}) &= \mathcal{M}_0(\hat{\mathbf{p}}_j) \cap \mathcal{M}_0\left(\sum_{j \in \mathbb{H}} (\hat{\mathbf{p}}_0 \cdot (\hat{\mathbf{r}}_j + \hat{\mathbf{r}}'_j) + \hat{\mathbf{p}}_j \cdot (\hat{\mathbf{r}}_0^j + \hat{\mathbf{r}}'_j))\right) \\
&\stackrel{\text{Lemma 4}}{=} \mathcal{M}_0(\hat{\mathbf{p}}_j) \cap \left(\bigcap_{j \in \mathbb{H}} \mathcal{M}_0((\hat{\mathbf{p}}_0 \cdot (\hat{\mathbf{r}}_j + \hat{\mathbf{r}}'_j) + \hat{\mathbf{p}}_j \cdot (\hat{\mathbf{r}}_0^j + \hat{\mathbf{r}}'_j)))\right) \\
&\stackrel{\text{Lemma 6}}{=} \mathcal{M}_0(\hat{\mathbf{p}}_j) \cap \left(\bigcap_{j \in \mathbb{H}} \mathcal{M}_0(\hat{\mathbf{p}}_0) \cap \mathcal{M}_0(\hat{\mathbf{p}}_j)\right) \\
&= \hat{S}'_\cap
\end{aligned}$$

Once the intermediate intersection is computed, the adversary can only add an update polynomial $\hat{\mathbf{p}}_{\text{upt}}$ to get the final intersection polynomial \mathbf{p}_\cap^* . It remains to show that this final intersection does not include any points that were not already in the intermediate intersection for any of the parties' polynomials $\hat{\mathbf{p}}_j$.

For this, we consider the intersection of every honest party's (unknown) input \mathbf{p}_j with the intersection. It has to hold that $\hat{S}'_\cap \subseteq \hat{S}_\cap$ for all P_j except with negligible probability. Here we require that $\Pr[x \leftarrow \mathcal{A}(\hat{\mathbf{p}}_{\text{int}}), \text{ s.t. } \hat{\mathbf{p}}_j(x) = 0 \wedge \hat{\mathbf{p}}_{\text{int}}(x) \neq 0] \leq \text{neg}!(|\mathbb{F}|)$, i.e. the adversary can only

guess an element of P_j 's input set.

$$\begin{aligned}
\mathcal{M}_0(\hat{\mathbf{p}}_j) \cap \mathcal{M}_0(\mathbf{p}_\cap^*) &= \mathcal{M}_0(\hat{\mathbf{p}}_j) \cap (\mathcal{M}_0(\hat{\mathbf{p}}_{\text{int}} + \hat{\mathbf{p}}_{\text{upt}})) \\
&\stackrel{\text{Lemma 4}}{=} \mathcal{M}_0(\hat{\mathbf{p}}_j) \cap (\mathcal{M}_0(\hat{\mathbf{p}}_{\text{int}}) \cap \mathcal{M}_0(\hat{\mathbf{p}}_{\text{upt}})) \\
&= \mathcal{M}_0(\hat{\mathbf{p}}_j) \cap (\hat{S}_\cap \cap \mathcal{M}_0(\hat{\mathbf{p}}_{\text{upt}})) \\
&\subseteq \mathcal{M}_0(\hat{\mathbf{p}}_j) \cap \hat{S}_\cap = \hat{S}_\cap
\end{aligned}$$

Therefore, $\hat{S}'_\cap \subseteq \hat{S}_\cap$, and the output in both hybrids is identical.

Regarding (b), we make the following observations. Since \mathcal{S}_4 sends $\hat{\mathbf{s}}'_j = \hat{\mathbf{s}}_j - \mathbf{u}_j + \sum_{i \neq j} \mathbf{v}_{ij}$, the value $\hat{\mathbf{s}}'_j$ is uniformly random over the choice of \mathbf{u}_j (and over $\sum \mathbf{v}_{ij}$, if $t \leq n - 2$). Therefore, the simulation of $\hat{\mathbf{s}}'_j$ is identically distributed to Hybrid 4.

Similarly, we have:

$$\sum_{j \in \mathbf{H}} (\hat{\mathbf{s}}'_j + \hat{\mathbf{s}}_0^j) = \sum_{j \in \mathbf{H}} (\hat{\mathbf{p}}_0 \cdot (\hat{\mathbf{r}}_j + \hat{\mathbf{r}}'_0) + \hat{\mathbf{p}}_j \cdot (\hat{\mathbf{r}}_0^j + \hat{\mathbf{r}}'_j)) \quad [+ \sum_{i \in \mathbf{A}, j \in \mathbf{H}} \mathbf{v}_{ij}]$$

We can ignore the \mathbf{v}_{ij} values, since these are known to \mathcal{A} . The sum is uniform over the choice of the $\hat{\mathbf{r}}_j, \hat{\mathbf{r}}'_j$ apart from the points $\gamma \in \hat{S}_\cap$ (since \mathcal{F}_{OPA} guarantees that $\hat{p}_0 \neq 0$) and therefore identically distributed to Hybrid 5, since the extraction is correct. \square

P_0 is honest: Consider the simulator in Figure 15.

Simulator $\mathcal{S}_{\bar{P}_0}$

Let $\mathbf{A} = \{i | P_i \text{ is malicious}\}$ denote the index set of corrupted parties, where $|\mathbf{A}| = t \leq n - 1$. Further let \mathbf{H} denote the index set of honest parties.

1. Simulate the setup and obtain all \mathbf{v}_{ij}^* for $i \in \mathbf{A}$ and $j \in \mathbf{H}$.
2. Extract the inputs $(\hat{\mathbf{p}}_i, \hat{\mathbf{r}}_i, \hat{\mathbf{u}}_i)$ for all $i \in \mathbf{A}$ by simulating \mathcal{F}_{OPA} .
3. Set $\hat{\mathbf{p}}_{\mathbf{A}} = \sum_{i \in \mathbf{A}} \hat{\mathbf{p}}_i$, and find the roots $\hat{\gamma}_1, \dots, \hat{\gamma}_{2m}$ of $\hat{\mathbf{p}}_{\mathbf{A}}$ and thereby the set $\hat{S}_{\mathbf{A}} = \{\hat{\gamma}_1, \dots, \hat{\gamma}_{2m}\}$.
4. Send $(\text{input}, P_i, \hat{S}_{\mathbf{A}})$ to $\mathcal{F}_{\text{MPSI}}^*$ for all parties $i \in \mathbf{A}$.
5. Upon receiving $(\text{output}, \hat{S}_\cap)$ from $\mathcal{F}_{\text{MPSI}}^*$, pick $n - t$ random degree m polynomials $\hat{\mathbf{p}}_j$ such that $\hat{\mathbf{p}}_j(\gamma) = 0$ for all $\gamma \in \hat{S}_\cap, j \in \mathbf{H}$.
6. Use the $\hat{\mathbf{p}}_j$ as input for each instance of \mathcal{F}_{OPA} together with random polynomials $\hat{\mathbf{r}}_0^i$ and $\hat{\mathbf{u}}_0^i$ for $i \in \mathbf{A}$, i.e. keep $\hat{\mathbf{s}}_0^i = \hat{\mathbf{p}}_0 \cdot \hat{\mathbf{r}}_i + \hat{\mathbf{u}}_i$ and send $\hat{\mathbf{s}}_i = \hat{\mathbf{p}}_i \cdot \hat{\mathbf{r}}_0^i + \hat{\mathbf{u}}_0^i$ to \mathcal{A} .
7. Simulate the rest according to Π_{MPSI} , but abort in Step 5, if after setting
$$\hat{\mathbf{p}}_{\text{test}} = \sum_{i \in \mathbf{A}} (\hat{\mathbf{s}}_i^* + \hat{\mathbf{u}}_i - \hat{\mathbf{u}}_0^i - \sum_{j \in \mathbf{H}} \mathbf{v}_{ij}^*) - \sum_{i \in \mathbf{A}} \hat{\mathbf{p}}_i \cdot \hat{\mathbf{r}}_0^i = \sum_{i \in \mathbf{A}} \hat{\mathbf{p}}_i \cdot \hat{\mathbf{r}}_i^*,$$
 $\alpha_i \neq \hat{\mathbf{p}}_i(x), \beta_i \neq \hat{\mathbf{r}}_i(x)$ or $\sum_{i \in \mathbf{A}} \alpha_i \delta_i \neq \hat{\mathbf{p}}_{\text{test}}(x)$ for any $i \in \mathbf{A}$ even if the check passes otherwise.

Figure 15: Simulator $\mathcal{S}_{\bar{P}_0}$ for $P_0 \notin \mathbf{A}$.

We show the indistinguishability of the simulation and the real protocol through the following hybrid games.

Hybrid 0: $\text{Real}_{\Pi_{\text{MPSI}}}^{\mathcal{A}}$.

Hybrid 1: Identical to Hybrid 0, except that \mathcal{S}_1 simulates \mathcal{F}_{OPA} , learns all inputs and aborts if $\alpha_i \neq \hat{\mathbf{p}}_i(x)$ or $\beta_i \neq \hat{\mathbf{r}}_i(x)$ for any $i \in \mathbf{A}$, but the check is passed.

Hybrid 2: Identical to Hybrid 1, except that \mathcal{S}_2 aborts according to Step 7 in Figure 15.

Hybrid 3: Identical to Hybrid 2, except that \mathcal{S}_3 generates the inputs $\hat{\mathbf{s}}_0^j, \hat{\mathbf{s}}_j$ according to Step 6 in Figure 15. This corresponds to $\text{Ideal}_{\mathcal{F}_{\text{MPSI}}}^{\mathcal{S}_{\hat{\mathbf{P}}_0}}$.

We prove the computational indistinguishability of Hybrids 0 and 1 in Lemma 9.4. This does not rule out that \mathcal{A} adds a masking polynomial, thereby leading to an incorrectly extracted input. Lemma 9.5 takes care of this problem by showing the indistinguishability of Hybrids 1 and 2. From this point on \mathcal{A} 's input is correctly defined, which allows to show that Hybrid 2 and Hybrid 3 are statistically close. Let us briefly argue that setting $\hat{\mathbf{p}}_{\mathbf{A}} = \sum_{i \in \mathbf{A}} \hat{\mathbf{p}}_i$ yields the correct intersection. With overwhelming probability over $\hat{\mathbf{p}}_j$ it holds that

$$\begin{aligned} \mathcal{M}_0(\hat{\mathbf{p}}_{\mathbf{A}}) &= \mathcal{M}_0\left(\sum_{i \in \mathbf{A}} \hat{\mathbf{p}}_i\right) \\ &\stackrel{\text{Lemma 4}}{=} \bigcap_{i \in \mathbf{A}} (\mathcal{M}_0(\hat{\mathbf{p}}_i)) \\ &= \bigcap_{i \in \mathbf{A}} S_i. \end{aligned}$$

Then, independent of \mathcal{A} 's inputs, $\hat{\mathbf{p}}_{\cap}$ returns the correct intersection \hat{S}_{\cap} for all $i \in \mathbf{A}$ except with negligible probability.

$$\begin{aligned} \mathcal{M}_0(\hat{\mathbf{p}}_i) \cap \mathcal{M}_0(\hat{\mathbf{p}}_{\cap}) &= \mathcal{M}_0(\hat{\mathbf{p}}_i) \cap \mathcal{M}_0\left(\sum_{l=1}^{n-1} (\hat{\mathbf{p}}_0 \cdot (\hat{\mathbf{r}}_l + \hat{\mathbf{r}}'_0) + \hat{\mathbf{p}}_l \cdot (\hat{\mathbf{r}}_0^l + \hat{\mathbf{r}}_l'))\right) \\ &\stackrel{\text{Lemma 4}}{=} \mathcal{M}_0(\hat{\mathbf{p}}_i) \cap \left(\bigcap_{l=1}^{n-1} \mathcal{M}_0((\hat{\mathbf{p}}_0 \cdot (\hat{\mathbf{r}}_l + \hat{\mathbf{r}}'_0) + \hat{\mathbf{p}}_l \cdot (\hat{\mathbf{r}}_0^l + \hat{\mathbf{r}}_l'))\right) \\ &\stackrel{\text{Lemma 6}}{=} \mathcal{M}_0(\hat{\mathbf{p}}_i) \cap \left(\bigcap_{l=1}^{n-1} \mathcal{M}_0(\hat{\mathbf{p}}_0) \cap \mathcal{M}_0(\hat{\mathbf{p}}_l)\right) \\ &= \hat{S}_{\cap} \end{aligned}$$

The polynomial $\hat{\mathbf{s}}_i = \hat{\mathbf{p}}_0 \cdot \hat{\mathbf{r}}_i + \hat{\mathbf{u}}_0^i$ and in particular its roots are uniformly distributed over the choice of $\hat{\mathbf{u}}_0^i$ from \mathcal{A} 's view. Replacing $\hat{\mathbf{p}}_0$ with $\hat{\mathbf{p}}_0'$ does not change this. Since we established that the extracted inputs are correct, the intersection polynomial $\hat{\mathbf{p}}_{\cap}$ contains exactly the set of the intersection as its roots, and is otherwise random over the choice of $\hat{\mathbf{r}}_0^i, \hat{\mathbf{r}}_0'$. Thus, the output of the simulation is identical to the output of the ideal functionality.

As a result we get that for all PPT environments \mathcal{Z} ,

$$\text{Real}_{\Pi_{\text{MPSI}}}^{\mathcal{A}}(\mathcal{Z}) \approx_c \text{Ideal}_{\mathcal{F}_{\text{MPSI}}}^{\mathcal{S}_{\hat{P}_0}}(\mathcal{Z}).$$

Lemma 9.4. *Assume that NMCOM is a bounded-concurrent non-malleable commitment scheme against synchronizing adversaries. Then Hybrid 0 and Hybrid 1 are computationally indistinguishable.*

Proof. The only difference between Hybrid 0 and Hybrid 1 lies in the fact that \mathcal{S}_1 aborts if the extracted $\hat{\mathbf{p}}_i$ evaluated on x does not match the check value α_i , but the check is still passed. Our proof follows along the lines of the proof of Lemma 9.1, with some small modifications.

In order for \mathcal{Z} to distinguish both hybrids, it has to be able to produce a value $\alpha_i^* \neq \hat{\mathbf{p}}_i$ and pass the check with non-negligible probability ϵ . W.l.o.g. it is sufficient if α_i^* is incorrect for only one $\hat{\mathbf{p}}_i$. We show that such a \mathcal{Z} breaks the non-malleability property of NMCOM.

Let σ denote the outcome of the check. If \mathcal{A} is honest, i.e. $\alpha_i = \hat{\mathbf{p}}_i(x)$, $\beta_i = \hat{\mathbf{r}}_i(x)$ and $\delta_i = \hat{\mathbf{r}}_i'(x)$, then

$$\sigma = \sum_{i=0}^n (\alpha_0(\beta_i + \delta_0) + \alpha_i(\beta_0^i + \delta_i)) - \mathbf{p}_{\cap}(x) = 0, \quad (2)$$

where

$$\mathbf{p}_{\cap} = \sum_{i \in \mathcal{A}} (\mathbf{s}_i + \mathbf{s}_0^i) + \sum_{j \in \mathcal{H}} (\mathbf{s}_j + \mathbf{s}_0^j).$$

Since \mathcal{A} does not control the inputs of the honest parties $j \in \mathcal{H}$, in order to pass the check it is sufficient for \mathcal{A} to satisfy the following simplification of Equation (2).

$$\sigma' = \sum_{i \in \mathcal{A}} (\alpha_0(\beta_i + \delta_0) + \alpha_i(\beta_0^i + \delta_i)) - \sum_{i \in \mathcal{A}} (\mathbf{s}_i(x) + \mathbf{s}_0^i(x)) = \text{const}$$

Here const is a fixed constant known to \mathcal{A} (0 if \mathcal{A} is honest) determined by setting the inputs for $\mathbf{s}_i, \mathbf{s}_0^i$ for $i \in \mathcal{A}$ accordingly. But if $\alpha_i^* \neq \hat{\mathbf{p}}_i(x)$, i.e. $\alpha_i^* = \alpha_i + e$, then we get that

$$\begin{aligned} \sigma' &= \sum_{i \in \mathcal{A}} (\alpha_0(\beta_i + \delta_0) + (\alpha_i + e)(\beta_0^i + \delta_i)) - \sum_{i \in \mathcal{A}} (\mathbf{s}_i(x) + \mathbf{s}_0^i(x)) \\ &= \sum_{i \in \mathcal{A}} (\alpha_0(\beta_i + \delta_0) + \alpha_i(\beta_0^i + \delta_i)) - \sum_{i \in \mathcal{A}} (\mathbf{s}_i(x) + \mathbf{s}_0^i(x)) + e(\beta_0^i + \delta_i) \\ &= e(\beta_0^i + \delta_i) \neq \text{const} \end{aligned}$$

Similarly for $\beta_i \neq \hat{\mathbf{r}}_i(x)$ for any $i \in \mathcal{A}$. Thus, except for the case of $\alpha_i^* = \alpha_i + e/\beta_0^i$, the check will fail for $\alpha_i^* \neq \hat{\mathbf{p}}_i(x)$. But since we assumed that \mathcal{A} passes the check with non-negligible probability, and NMCOM is statistically binding, \mathcal{A} has to produce a valid commitment to $\tilde{\alpha}_i = \alpha_i + e/(\beta_0^i + \delta_i)$ with the same probability.

Consider the two views

$$\text{Real} = \{\hat{\mathbf{p}}_{\cap}, \{\text{com}_j\}_{j \in \mathcal{H}}\}, \quad \text{Rand} = \{\hat{\mathbf{p}}_{\cap}, \{\widehat{\text{com}}_j\}_{j \in \mathcal{H}}\},$$

where $\text{com}_j \leftarrow \text{NMCOM.Commit}(\alpha_j, \beta_j, \delta_j)$ and $\widehat{\text{com}}_j \leftarrow \text{NMCOM.Commit}(0)$. Real corresponds to a real protocol view of \mathcal{A} before committing itself. Obviously, $\text{Real} \approx_c \text{Rand}$ if NMCOM is non-malleable. However, we will argue that \mathcal{A} cannot output a valid commitment on $\tilde{\alpha}_i$ except with negligible probability, i.e.

$$\Pr[(\text{com}_i^*, \text{unv}_i^*, (\tilde{\alpha}_i, \tilde{\beta}_i, \tilde{\delta}_i) \leftarrow \mathcal{A}(\text{Rand}) \wedge \text{valid}] \leq \text{negl}(\kappa),$$

where valid is the event that $\text{NMCOM.Open}(\text{com}_i^*, \text{unv}_i^*, (\tilde{\alpha}_i, \tilde{\beta}_i, \tilde{\delta}_i)) = 1$. We first observe that $\hat{\mathbf{p}}_0$ and $\hat{\mathbf{r}}_0^i$ for $i \in \mathbf{A}$ cannot be obtained by \mathcal{A} via $\hat{\mathbf{p}}_\cap$:

$$\begin{aligned} \hat{\mathbf{p}}_\cap &= \sum_{i \in \mathbf{A}} (\mathbf{s}_i + \mathbf{s}_0^i) + \sum_{j \in \mathbf{H}} (\mathbf{s}_j + \mathbf{s}_0^j) \\ &= \sum_{i \in \mathbf{A}} (\hat{\mathbf{p}}_0 \cdot (\hat{\mathbf{r}}_i + \hat{\mathbf{r}}_0') + \hat{\mathbf{p}}_i \cdot (\hat{\mathbf{r}}_0^i + \hat{\mathbf{r}}_i')) + \sum_{j \in \mathbf{H}} (\hat{\mathbf{p}}_0 \cdot (\hat{\mathbf{r}}_j + \hat{\mathbf{r}}_0') + \hat{\mathbf{p}}_j \cdot (\hat{\mathbf{r}}_0^j + \hat{\mathbf{r}}_j')) \end{aligned}$$

Thus, having only one equation, but at least two unknowns, it is information-theoretically impossible to reconstruct $\hat{\mathbf{p}}_0$, $\hat{\mathbf{r}}_0^i$ or $\hat{\mathbf{r}}_0'$ (and therefore $\tilde{\alpha}_i$) except with negligible probability. Further, the commitments com_j contain the value 0 and are therefore independent of $\hat{\mathbf{p}}_0$, $\hat{\mathbf{r}}_0^i$ and $\hat{\mathbf{r}}_0'$. Thus, the probability that \mathcal{A} obtains a commitment on $\tilde{\alpha}_i$ is negligible.

But since $\text{Real} \approx_c \text{Rand}$, we also get that

$$\Pr[(\text{com}_i^*, \text{unv}_i^*, (\tilde{\alpha}_i, \tilde{\beta}_i, \tilde{\delta}_i) \leftarrow \mathcal{A}(\text{Real}) \wedge \text{valid}] \leq \text{negl}(\kappa),$$

which contradicts our assumption that \mathcal{A} produces the commitment with non-negligible probability ϵ .

In conclusion, Hybrid 0 and Hybrid 1 are computationally indistinguishable. \square

Lemma 9.5. *Assume that Π_{CT} provides a uniformly random x with computational security. Then Hybrid 1 and Hybrid 2 are computationally indistinguishable.*

Proof. The only difference between Hybrid 1 and Hybrid 2 lies in the fact that \mathcal{S}_2 aborts if

$$\mathbf{p}_{\text{test}}(x) \neq \sum_{i \in \mathbf{A}} \alpha_i \delta_i,$$

while \mathcal{S}_1 does not.

If \mathcal{Z} wants to distinguish Hybrid 1 and Hybrid 2, it thus has to provide a value \mathbf{s}_i^{*} for which the above equation does not hold, but at the same time the check in Step 5 has to pass. In other words, for at least one malicious party P_i , \mathcal{A} either does not use the same $\hat{\mathbf{u}}_i$ as input in \mathcal{F}_{OPA} and for $\hat{\mathbf{s}}_i^{*}$, or the same \mathbf{v}_{ij} , leaving $\hat{\mathbf{p}}_i \cdot (\hat{\mathbf{r}}_0^i + \hat{\mathbf{r}}_i') + \mathbf{f}$, where \mathbf{f} is a non-zero polynomial of degree $2m$ and in particular not of the form $\hat{\mathbf{p}}_i \cdot \mathbf{q}$ for some degree m polynomial \mathbf{q} . Intuitively, this means that if \mathcal{A} provides such an input and passes the check, the extracted input is incorrect (because \mathcal{S} only looks at the $\hat{\mathbf{p}}_i$ values, and \mathbf{f} changes the intersection polynomial non-trivially).

Lemma 9.4 establishes that if $\alpha_i \neq \hat{\mathbf{p}}_i(x)$ or $\beta_i \neq \hat{\mathbf{r}}_i(x)$, then the check fails. It follows that the check can only be passed if $\mathbf{f}(x) = 0$, since otherwise $\alpha_i = \hat{\mathbf{p}}_i(x) + \mathbf{f}(x)$.

Assume that \mathcal{Z} succeeds in this endeavour with polynomial probability ϵ . Then we can construct from \mathcal{Z} an adversary \mathcal{B} that predicts the outcome of Π_{CT} with polynomial probability.

First note that \mathbf{f} is of degree $2m$ and the probability that $\mathbf{f}(x) = 0$ in a random point x over \mathbb{F} is $2m/|\mathbb{F}|$, which is negligible in our case. Our adversary \mathcal{B} proceeds as follows. It simulates the protocol for \mathcal{Z} according to \mathcal{S}_1 up to the point where \mathcal{S}_1 learns the polynomials $\hat{\mathbf{p}}_i, \hat{\mathbf{r}}_i$ and $\hat{\mathbf{u}}_i$. After receiving $\hat{\mathbf{s}}_i^{/*}$ for $i \in \mathbf{A}$, \mathcal{B} sets

$$\mathbf{f} = \sum_{i \in \mathbf{A}} (\mathbf{s}_i^{/*} + \hat{\mathbf{u}}_i - \hat{\mathbf{u}}_0^i - \sum_{j \in \mathbf{H}} \mathbf{v}_{ij}^*) - \sum_{i \in \mathbf{A}} \hat{\mathbf{p}}_i \cdot \hat{\mathbf{r}}_0^i$$

and computes the roots $\gamma_1, \dots, \gamma_{2m}$ of \mathbf{f} . One of these roots has to be the biased point x , otherwise $\mathbf{f}(x) \neq 0$ and the check in Π_{MPSI} fails. \mathcal{B} picks a random $j \in [2m]$ and predicts the output of the coin-flip as γ_j . Thus, \mathcal{B} predicts the outcome of the coin-toss correctly with probability $\epsilon/2m$, which is polynomial. This contradicts the security of Π_{CT} .

In conclusion, Hybrid 1 and Hybrid 2 are computationally indistinguishable. \square

Efficiency. The setup, i.e. the distribution of seeds, has communication complexity $O(n^2\kappa)$.

The oblivious addition of the polynomials has communication overhead of $O(nm\kappa)$. The check phase first requires a multi-party coin-toss.

We propose to use a standard Blum coin-toss and generalise it as follows to n parties. Let Π_{CT} proceed as follows. Every party picks a random field element x_i . It commits to the field element using a concurrent non-malleable commitment (similar to the check in Π_{MPSI}). Then, every party broadcasts x_i and verifies that the commitment is correct. If that holds, the output is defined as $x_{\text{CT}} = \sum x_i$, otherwise party P_i aborts. The communication complexity of Π_{CT} is in $O(n^2 \cdot c_{\text{com}})$, where c_{com} is the communication overhead for the commitment used.

The check itself also requires sending $O(n^2)$ non-malleable commitments. The most efficient (bounded-)concurrent non-malleable commitment that we are aware of is due to Goyal et al. [GRRV14] (concurrency was shown in [COSV17]). This commitment has communication complexity $O(\kappa^2)$ in our setting, i.e. the check and coin-toss would require communication in $O(n^2\kappa^2)$.

Therefore, in order to achieve an asymptotically optimal construction, we opt to use an OLE-based UC-secure commitment from the supplementary material 2.2 instead of the non-malleable one. UC-security implies concurrent non-malleability in the strongest form, so that this step does not jeopardize the security of the protocol. This commitment has a constant rate, i.e. the communication complexity of the check is reduced to $O(n^2\kappa)$. Combining the above observations, Π_{MPSI} has communication complexity $O((n^2 + nm)\kappa)$ in the \mathcal{F}_{OLE} -hybrid model.

For concrete instantiations of \mathcal{F}_{OLE} , the OLE protocol of Ghosh et al. [GNN17] has a constant communication overhead per OLE. In summary, the complete protocol has communication complexity $O((n^2 + nm)\kappa)$, which is asymptotically optimal for $m \geq n$.

Similar to the two-party case, the computational cost is dominated by the cost of polynomial interpolation. In particular, the central party has to run the two-party protocol n times, which leads to a computational overhead of $O(nm \log m)$. The other parties basically have the same computational overhead as in the two-party case. \square

7 Threshold-(M)PSI

In this section we will show a simple modification of our results from the previous sections. Specifically, we construct threshold set intersection protocols with only a small constant overhead over our standard (M)PSI protocol.

In Figure 16 we formally define the ideal functionality $\mathcal{F}_{\text{MPSI}}^\ell$ for multi-party threshold PSI. The only difference from the standard MPSI functionality lies in the additional check whether the size of the intersection exceeds a threshold ℓ .




Figure 16: Ideal functionality $\mathcal{F}_{\text{MPSI}}^\ell$ for multi-party threshold PSI.

7.1 Threshold Multi-party PSI from OLE

We simply have to modify the inputs of the parties in comparison to the standard multi-party PSI protocol. Instead of simply setting \mathbf{p}_i such that $\mathbf{p}_i(\gamma_j) = 0$ for all $\gamma_j \in S_i$, we set $\mathbf{p}_i(\gamma_j) = 1$. Further, for each of the random polynomials $\mathbf{r}_i, \mathbf{r}'_i$ we set $\mathbf{r}_i(\gamma_j) = \rho_j$ and $\mathbf{r}'_i(\gamma_j) = \rho'_j$, where $\rho_1, \dots, \rho_n, \rho'_1, \dots, \rho'_n$ are the shares of two robust (ℓ, n) -secret sharings of random values s_i^0 and s_i^1 , respectively. After executing the protocol with these inputs, the parties obtain the correct intersection if each party reconstructs the correct value $r_\cap = \sum_{i=1}^n (s_i^0 + s_i^1)$.

Security of this protocol follows from the fact that the shares ρ_j are uniformly random, so \mathbf{r}_i and \mathbf{r}'_i are distributed exactly as in Π_{MPSI} , unless a party obtains more than ℓ shares. If more than ℓ shares are obtained, a malicious party learns at most the sum of the two random values s_i^0 and s_i^1 , which reveals nothing about the individual polynomials \mathbf{r}_i and \mathbf{r}'_i . Also, choosing \mathbf{p}_i as described above is a simple input substitution for Π_{MPSI} , so this does not lead to any problems.

In order to prevent rushing adversaries from cheating, the parties additionally have to commit to their value r_\cap , and then broadcast this commitment before checking the consistency. If this check passes, it means that the intersection contains more than t elements, and the robustness of the secret sharing scheme ensures that the randomized shares do not break the reconstruction. Importantly, if the parties do not obtain enough shares to reconstruct r_\cap , the intersection remains completely hidden, since the shares are uniformly distributed.

Lemma 10. *The protocol Π_{MPSI}^ℓ UC-realizes $\mathcal{F}_{\text{MPSI}}^\ell$ in the \mathcal{F}_{OLE} -hybrid model.*

Proof. Let us briefly sketch our proof strategy. Since Π_{MPSI}^ℓ basically runs an instance of Π_{MPSI} with modified inputs and a slightly modified output phase, we will write a simulator for Π_{MPSI} which translates any attack on Π_{MPSI}^ℓ into an attack on Π_{MPSI} . Again, we have to distinguish between a malicious P_0 and an honest one; however, the proof strategy in both cases is identical, so we only show the first (and arguably more complicated) case.

We show the indistinguishability of the simulation and the real protocol through the following hybrid games. In the following, let \mathcal{A} denote the dummy adversary controlled by \mathcal{Z} .

Hybrid 0: $\text{Real}_{\Pi_{\text{MPSI}}}^{\mathcal{A}}$.

$$\Pi_{\text{MPSI}}^\ell$$

Let $m = \max_i \{|\mathcal{S}_i|\} + 1$ and RSS a robust (ℓ, n) -secret sharing.

Setup Identical to Π_{MPSI} .

Share Computation

1. P_0 (Input S_0): Instantiate the inputs as follows.

- Compute a polynomial \mathbf{p}_0 of degree m s.t. $\mathbf{p}_0(\gamma_j) = 1$ for all $\gamma_j \in S_0$.
- Pick two random values $s_0^0, s_0^1 \in \mathbb{F}$ and compute two secret sharings $(\rho_1, \dots, \rho_m) \leftarrow \text{RSS.Share}(s_0^0)$ and $(\rho'_1, \dots, \rho'_m) \leftarrow \text{RSS.Share}(s_0^1)$.
- Generate n random polynomials $\mathbf{r}_0^i \in \mathbb{F}[X], i \in \{1, \dots, n-1\}$ and $\mathbf{r}'_0 \in \mathbb{F}[X]$ of degree m each s.t. $\mathbf{r}_0^i(\gamma_j) = \rho_j$ and $\mathbf{r}'_0(\gamma_j) = \rho'_j$ for all $\gamma_j \in S_0$.
- Generate $n-1$ random polynomials $\mathbf{u}_0^i \in \mathbb{F}[X], i \in \{1, \dots, n-1\}$ of degree $2m$.

Proceed with Π_{MPSI} using the above inputs.

2. P_i (Input S_i): Instantiate the inputs as follows.

- Compute a polynomial \mathbf{p}_i of degree m s.t. $\mathbf{p}_i(\gamma_j) = 1$ for all $\gamma_j \in S_i$.
- Pick two random values $s_i^0, s_i^1 \in \mathbb{F}$ and compute two secret sharings $(\rho_1, \dots, \rho_m) \leftarrow \text{RSS.Share}(s_i^0)$ and $(\rho'_1, \dots, \rho'_m) \leftarrow \text{RSS.Share}(s_i^1)$.
- Generate two polynomials $\mathbf{r}_i, \mathbf{r}'_i \in \mathbb{F}[X]$ of degree m s.t. $\mathbf{r}_i(\gamma_j) = \rho_j$ and $\mathbf{r}'_i(\gamma_j) = \rho'_j$ for all $\gamma_j \in S_i$.
- Generate $\mathbf{u}_i \in \mathbb{F}[X]$ uniformly of degree $2m$.

Proceed with Π_{MPSI} using the above inputs.

Output Aggregation and Verification

3. Run the output aggregation and verification phase of Π_{MPSI} and obtain \mathbf{p}_\cap .

4. *All parties:*

- Compute $r_\cap^i = \text{RSS.Recon}(\mathbf{p}_\cap(\gamma_1), \dots, \mathbf{p}_\cap(\gamma_m))$.
- Compute $(\text{com}_r^i, \text{unv}_r^i) = \text{NMCOM.Commit}(r_\cap^i)$ and broadcast com_r^i .
- Once all commitments are received, broadcast $(r_\cap^i, \text{unv}_r^i)$ and abort if $\text{NMCOM.Open}(\text{com}_r^i, \text{unv}_r^i, r_\cap^i) \neq 1$ or $r_\cap^i \neq r_\cap^j$ for any $i, j \in [n]$.
- For all $j \in [m]$, if $\mathbf{p}_\cap(\gamma_j)$ is a valid share of r_\cap , add γ_j to S_\cap .

Figure 17: Protocol Π_{MPSI}^ℓ for threshold set intersection in the \mathcal{F}_{OLE} -hybrid model.

Hybrid 1: Identical to Hybrid 0, except that \mathcal{S}_1 aborts if the check in Step 4 of Π_{MPSI}^ℓ passes, even though $\hat{S}_\cap = \emptyset$.

Hybrid 2: Identical to Hybrid 1, except that \mathcal{S}_2 substitutes all inputs by running \mathcal{S}_{P_0} . This

Simulator $\mathcal{S}_{P_0}^\ell$

Let $\mathbf{A} = \{i | P_i \text{ is malicious}\}$ denote the index set of corrupted parties, where $|\mathbf{A}| = t \leq n - 1$. Further let \mathbf{H} denote the index set of honest parties.

1. Run \mathcal{S}_{P_0} of Π_{MPSI} with the only modification that it extracts the roots of the polynomial $\mathbf{p}_{\mathbf{A}} - 1$ (i.e. it returns all γ_j such that $\mathbf{p}_{\mathbf{A}}(\gamma_j) = 1$). Forward all messages from \mathcal{A} to \mathcal{S}_{P_0} .
2. Forward $(\text{input}, P_i, \hat{S}_{\mathbf{A}})$ from \mathcal{S}_{P_0} to $\mathcal{F}_{\text{MPSI}}^\ell$ for all parties $i \in \mathbf{A}$.
3. Upon receiving $(\text{output}, \hat{S}_{\cap})$ from $\mathcal{F}_{\text{MPSI}}^\ell$, forward this message to \mathcal{S}_{P_0} and continue the simulation (with the exception that the $\hat{\mathbf{p}}_j(\gamma_j) = 1$ for $\gamma_j \in \hat{S}_{\cap}$ instead of 0). Forward all messages from \mathcal{A} to \mathcal{S}_{P_0} .
4. Upon receiving $\text{com}_r^{\mathbf{A}}$, run the reconstruction protocol of r_{\cap} according to Π_{MPSI}^ℓ . Abort if $\hat{S}_{\cap} = \emptyset$, but the check in Step 4 passes.

Figure 18: Simulator $\mathcal{S}_{P_0}^\ell$ for $P_0 \in \mathbf{A}$.

corresponds to $\text{Ideal}_{\mathcal{F}_{\text{MPSI}}^*}^{\mathcal{S}_{P_0}}$.

We only sketch the proof, because the argumentation follows along the lines of the previous proofs, with minor modifications. The first two hybrids are indistinguishable due to the non-malleability of NMCOM and the fact that the s_i^0, s_i^1 are chosen uniformly. Hybrid 1 and Hybrid 2 are also indistinguishable: all shares of RSS are uniformly distributed, exactly like the random polynomials in Π_{MPSI} . All messages except for the commitment check at the end are identical to Π_{MPSI} , so the simulation of $\mathcal{S}_{P_0}^\ell$ is also indistinguishable from the real protocol. Finally, the output is also identical: The previous step together with the correctness of the simulation of \mathcal{S}_{P_0} already establishes that \mathcal{A} cannot convince an honest party of an incorrect intersection, since \mathcal{S}_{P_0} always extracts the correct input. Therefore, the slightly modified simulation will return the correct output and both hybrids are indistinguishable. \square

Relaxing the Threshold. Π_{MPSI}^ℓ works only for $\ell = \frac{2}{3}$ when Shamir's scheme [Sha79] is used to instantiate the robust secret sharing scheme. However, the protocol can easily be modified to work with any threshold with communication overhead less than 2 over Π_{MPSI}^ℓ . The idea is to add dummy elements in input sets to adjust the threshold artificially. The strategy works with any other robust secret sharing scheme as well. We will sketch the solution for the two relevant cases.

$\ell < \frac{2}{3}$: One party announces $\frac{2}{3}n - \ell$ elements which all parties include in their set. This increases the overall size, and pushes the size of the intersection above the reconstruction threshold.

$\ell > \frac{2}{3}$: All parties pick $\frac{3}{2}\ell - n$ random elements from \mathbb{F} . This is only necessary if the threshold of the robust secret sharing scheme cannot be fixed arbitrarily for $\ell \geq \frac{2}{3}$

Note, that a malicious party cannot influence the outcome except by forcing an abort (by not adding the announced element to the intersection), which cannot be prevented in any case.

Acknowledgement We would like to thank Ivan Damgård and Claudio Orlandi for helpful discussions.

References

- [ADI⁺17] Benny Applebaum, Ivan Damgård, Yuval Ishai, Michael Nielsen, and Lior Zichron. Secure arithmetic computation with constant computational overhead. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part I*, volume 10401 of *LNCS*, pages 223–254. Springer, Heidelberg, August 2017.
- [BFO12] Eli Ben-Sasson, Serge Fehr, and Rafail Ostrovsky. Near-linear unconditionally-secure multiparty computation with a dishonest minority. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 663–680. Springer, Heidelberg, August 2012.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd FOCS*, pages 136–145. IEEE Computer Society Press, October 2001.
- [CDI05] Ronald Cramer, Ivan Damgård, and Yuval Ishai. Share conversion, pseudorandom secret-sharing and applications to secure computation. In Joe Kilian, editor, *TCC 2005*, volume 3378 of *LNCS*, pages 342–362. Springer, Heidelberg, February 2005.
- [CJS12] Jung Hee Cheon, Stanislaw Jarecki, and Jae Hong Seo. Multi-party privacy-preserving set intersection with quasi-linear complexity. *IEICE Transactions*, 95-A(8):1366–1378, 2012.
- [COSV17] Michele Ciampi, Rafail Ostrovsky, Luisa Siniscalchi, and Ivan Visconti. Four-round concurrent non-malleable commitments from one-way functions. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part II*, volume 10402 of *LNCS*, pages 127–157. Springer, Heidelberg, August 2017.
- [DCW13] Changyu Dong, Liqun Chen, and Zikai Wen. When private set intersection meets big data: an efficient and scalable protocol. In Ahmad-Reza Sadeghi, Virgil D. Gligor, and Moti Yung, editors, *ACM CCS 13*, pages 789–800. ACM Press, November 2013.
- [DKT10] Emiliano De Cristofaro, Jihye Kim, and Gene Tsudik. Linear-complexity private set intersection protocols secure in malicious model. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 213–231. Springer, Heidelberg, December 2010.
- [DMRY09] Dana Dachman-Soled, Tal Malkin, Mariana Raykova, and Moti Yung. Efficient robust private set intersection. In Michel Abdalla, David Pointcheval, Pierre-Alain Fouque, and Damien Vergnaud, editors, *ACNS 09*, volume 5536 of *LNCS*, pages 125–142. Springer, Heidelberg, June 2009.
- [DT10] Emiliano De Cristofaro and Gene Tsudik. Practical private set intersection protocols with linear complexity. In Radu Sion, editor, *FC 2010*, volume 6052 of *LNCS*, pages 143–159. Springer, Heidelberg, January 2010.
- [FHNP16] Michael J. Freedman, Carmit Hazay, Kobbi Nissim, and Benny Pinkas. Efficient set intersection with simulation-based security. *Journal of Cryptology*, 29(1):115–155, January 2016.

- [FNP04] Michael J. Freedman, Kobbi Nissim, and Benny Pinkas. Efficient private matching and set intersection. In Christian Cachin and Jan Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 1–19. Springer, Heidelberg, May 2004.
- [GNN17] Satrajit Ghosh, Jesper Buus Nielsen, and Tobias Nilges. Maliciously secure oblivious linear function evaluation with constant overhead. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 629–659. Springer, Heidelberg, December 2017.
- [GRRV14] Vipul Goyal, Silas Richelson, Alon Rosen, and Margarita Vald. An algebraic approach to non-malleability. In *55th FOCS*, pages 41–50. IEEE Computer Society Press, October 2014.
- [Haz15] Carmit Hazay. Oblivious polynomial evaluation and secure set-intersection from algebraic PRFs. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015, Part II*, volume 9015 of *LNCS*, pages 90–120. Springer, Heidelberg, March 2015.
- [HEK12] Yan Huang, David Evans, and Jonathan Katz. Private set intersection: Are garbled circuits better than custom protocols? In *NDSS 2012*. The Internet Society, February 2012.
- [HFH99] Bernardo A. Huberman, Matt Franklin, and Tad Hogg. Enhancing privacy and trust in electronic communities. In *Proceedings of the 1st ACM Conference on Electronic Commerce*, EC '99, pages 78–86, 1999.
- [HN12] Carmit Hazay and Kobbi Nissim. Efficient set operations in the presence of malicious adversaries. *Journal of Cryptology*, 25(3):383–433, July 2012.
- [HOS17] Per A. Hallgren, Claudio Orlandi, and Andrei Sabelfeld. Privatepool: Privacy-preserving ridesharing. In *30th IEEE Computer Security Foundations Symposium, CSF 2017, Santa Barbara, CA, USA, August 21-25, 2017*, pages 276–291, 2017.
- [HV17] Carmit Hazay and Muthuramakrishnan Venkitasubramaniam. Scalable multi-party private set-intersection. In Serge Fehr, editor, *PKC 2017, Part I*, volume 10174 of *LNCS*, pages 175–203. Springer, Heidelberg, March 2017.
- [IPS09] Yuval Ishai, Manoj Prabhakaran, and Amit Sahai. Secure arithmetic computation with no honest majority. In Omer Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 294–314. Springer, Heidelberg, March 2009.
- [KKRT16] Vladimir Kolesnikov, Ranjit Kumaresan, Mike Rosulek, and Ni Trieu. Efficient batched oblivious PRF with applications to private set intersection. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 16*, pages 818–829. ACM Press, October 2016.
- [KMP⁺17] Vladimir Kolesnikov, Naor Matania, Benny Pinkas, Mike Rosulek, and Ni Trieu. Practical multi-party private set intersection from symmetric-key techniques. In Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu, editors, *ACM CCS 17*, pages 1257–1272. ACM Press, October / November 2017.

- [KMRS14] Seny Kamara, Payman Mohassel, Mariana Raykova, and Seyed Saeed Sadeghian. Scaling private set intersection to billion-element sets. In Nicolas Christin and Reihaneh Safavi-Naini, editors, *FC 2014*, volume 8437 of *LNCS*, pages 195–215. Springer, Heidelberg, March 2014.
- [KS05] Lea Kissner and Dawn Xiaodong Song. Privacy-preserving set operations. In Victor Shoup, editor, *CRYPTO 2005*, volume 3621 of *LNCS*, pages 241–257. Springer, Heidelberg, August 2005.
- [MM87] C. Meadows and D. Mutchler. Matching secrets in the absence of a continuously available trusted authority. *IEEE Transactions on Software Engineering*, SE-13(2):289–292, Feb 1987.
- [NPS99] Moni Naor, Benny Pinkas, and Reuban Sumner. Privacy preserving auctions and mechanism design. In *EC*, pages 129–139, 1999.
- [NTL⁺11] Arvind Narayanan, Narendran Thiagarajan, Mugdha Lakhani, Michael Hamburg, and Dan Boneh. Location privacy via private proximity testing. In *NDSS 2011*. The Internet Society, February 2011.
- [OOS17] Michele Orrù, Emanuela Orsini, and Peter Scholl. Actively secure 1-out-of-n OT extension with application to private set intersection. In *Topics in Cryptology - CT-RSA 2017 - The Cryptographers' Track at the RSA Conference 2017, San Francisco, CA, USA, February 14-17, 2017, Proceedings*, pages 381–396, 2017.
- [PSSZ15] Benny Pinkas, Thomas Schneider, Gil Segev, and Michael Zohner. Phasing: Private set intersection using permutation-based hashing. In *24th USENIX Security Symposium, USENIX Security 15, Washington, D.C., USA, August 12-14, 2015.*, pages 515–530, 2015.
- [PSZ14] Benny Pinkas, Thomas Schneider, and Michael Zohner. Faster private set intersection based on OT extension. Cryptology ePrint Archive, Report 2014/447, 2014. <http://eprint.iacr.org/2014/447>.
- [PSZ16] Benny Pinkas, Thomas Schneider, and Michael Zohner. Scalable private set intersection based on OT extension. Cryptology ePrint Archive, Report 2016/930, 2016. <http://eprint.iacr.org/2016/930>.
- [RR17] Peter Rindal and Mike Rosulek. Improved private set intersection against malicious adversaries. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 235–259. Springer, Heidelberg, May 2017.
- [Sha79] Adi Shamir. How to share a secret. *Commun. ACM*, 22(11):612–613, 1979.
- [Sha80] Adi Shamir. *On the power of commutativity in cryptography*, pages 582–595. Springer Berlin Heidelberg, Berlin, Heidelberg, 1980.

- [SS08] Yingpeng Sang and Hong Shen. Privacy preserving set intersection based on bilinear groups. In *Proceedings of the Thirty-first Australasian Conference on Computer Science - Volume 74*, ACSC '08, pages 47–54, 2008.